Report # DAAG53-76-C-0195

AUTOMATED IMAGE ENHANCEMENT TECHNIQUES

2ND GENERATION FLIR

D. H. Tack

P. Narendra

Interim Report
July 1, 1976 to December 1, 1976

Prepared for

U.S. ARMY ELECTRONICS COMMAND
AMSEL-NV-VI
NIGHT VISION LABORATORY
Fort Belvoir, Virginia 22060

HONEYWELL INC.
Systems and Research Center
2600 Ridgway Parkway
Minneapolis, Minnesota 55413

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOV'T ACCESSION NUMBER	3. RECIPIENT'S CATALOG NUMBER

Interim progress rept. no. 1, 1 Jul-1 Dec 76, on Phase 1,

4. TITLE (AND SUBTITLE)	5. TYPE OF REPORT/PERIOD COVERED
Automated Image Enhancement Techniques for 2nd Generation FLIR	Interim Progress Report July 1 to December 1, 1976
Second	6. PERFORMING ORG. REPORT NUMBER 76SRC57 *76-SRC/57*
7. AUTHOR(S) D. H. Tack P. Narendra	8. CONTRACT OR GRANT NUMBER(S) DAAG53-76-C-0195
9. PERFORMING ORGANIZATIONS NAME/ADDRESS Honeywell Inc., Systems and Research 2600 Ridgway Parkway, N.E. Minneapolis, Minnesota 55413	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME/ADDRESS U.S. Army Electronics Command Amsel-NV-VI, Night Vision Laboratory Fort Belvoir, Virginia 22060	12. REPORT DATE December 20, 1976 *20 Dec 76*
	13. NUMBER OF PAGES 71 *71 p.*
14. MONITORING AGENCY NAME/ADDRESS (IF DIFFERENT FROM CONT. OFF.)	15. SECURITY CLASSIFICATION (OF THIS REPORT) Unclassified
	15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (OF THIS REPORT)

Unlimited Distribution

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

17. DISTRIBUTION STATEMENT (OF THE ABSTRACT ENTERED IN BLOCK 20, IF DIFFERENT FROM REPORT)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (CONTINUE ON REVERSE SIDE IF NECESSARY AND IDENTIFY BY BLOCK NUMBER)

20. ABSTRACT (CONTINUE ON REVERSE SIDE IF NECESSARY AND IDENTIFY BY BLOCK NUMBER)

This report is the first interim progress report for Contract DAAG53-76-C-0195, Automated Image Enhancement Techniques for 2nd Generation FLIR. It describes results of Phase I studies of image processing applications to FLIR imaging systems from July 1, 1976 to December 1, 1976.

402 349

HD-168 REV 11/74

DD FORM 1473 1 JAN 73   EDITION OF 1 NOV 55 IS OBSOLETE

TABLE OF CONTENTS

3

LIST OF ILLUSTRATIONS

LIST OF ILLUSTRATIONS (Continued)

LIST OF ILLUSTRATIONS (Concluded)

SECTION I

DC RESTORATION

INTRODUCTION

This report is the first interim progress report for Contract DAAG53-76-C-0195, Automated Image Enhancement Techniques for 2nd Generation FLIR. It describes results of Phase I studies of image processing applications to FLIR imaging systems from July 1, 1976 to December 1, 1976.

The report is subdivided into five major headings corresponding to Phase I program tasks:

- DC Restoration

- Contrast Enhancement

- MRT Enhancement

- Resolution Restoration

- Statistical Analysis

DC restoration treats the problems of correcting or compensating for deficiencies in detectors and signal processing electronics which can result in harmful distortion of the IR image. Contrast enhancement deals with problems of transforming thermal images of wide dynamic range into the limited dynamic range of the display medium without losing essential information content. MRT enhancement involves spatial or temporal smoothing for noise

7

reduction. Resolution restoration treats detail recovery from imagery which has been blurred by the optics and electronics of the sensor. Lastly, statistical analysis of FLIR imagery is performed to provide quantifiable characterizations through shape, brightness and textural feature statistics.

DC RESTORATION

Direct versus AC Coupling

DC restoration traditionally refers to techniques for correcting imagery artifacts introduced by AC coupling detectors to their amplifying electronics. The reasons for AC coupling (usually with simple high pass RC networks) are:

- Background suppression for contrast rendition and efficient utilization of electronic dynamic range

- Removal of detector DC bias

- 1/f noise suppression

Artifacts created by AC coupling are transient effects (drooping across the scan line and undershooting on hot-cold transitions) and information loss about interchannel DC levels occurs. All channels have zero average signal.

The most common DC restore technique, which corrects the second defect, clamps the signals on all channels by imaging a common reference source prior to each scan. Each channel then measures scene radiance changes relative to the same reference level, and relative interchannel DC levels are preserved.

8

With this approach it is difficult to adapt the reference-source temperature
to the ambient-scene temperature. As a result, the dynamic range of the
electronics may be underutilized, and needless saturation at the hot or
cold end of the range will occur if the respective references are too hot or
too cold.

With the advent of monolithic and hybrid focal plane detector/CCD proces-
sors, it is uncertain whether or not the previous reasons for AC coupling
detector outputs are still valid.

Directly coupled PV IR/CCDs exhibit negligible 1/f noise (discussed later
in this section). This was one of the important reasons for AC coupling
when using PC detectors. Since a compromise between the deleterious
effects of 1/f noise and AC coupling transients are no longer needed, RC
time constants that are long relative to a line scan interval (RC > 1 sec for
100 channels and < 1% droop across a scan line) are now desired. On chip
capacitors up to about 100 pf are practical, but fabrication is difficult
when in series rather than shunt. Also, RC > 1 sec required R $> 10^{10}$ ohms,
which may be difficult to achieve as the input impedance to the CCD input
amplifier.

Adaptive background subtraction ideally ensures that the coldest part of the
scene is always at, or slightly above, the zero signal level. Prevailing
direct coupled PV IR/CCD techniques handle this at the CCD input stage by
variable gate thresholding which transfers only the charge above a background
level determined by the previous frame. A comparable adaptive approach
for a clamped, AC-coupled system would require injecting a scene dependent
charge directly on the coupling capacitor during the clamping interval,

9

rather than imaging a constant reference temperature. By so doing, however, the advantage of simple cancellation of detector DC bias through AC coupling would be lost. Rather, as in the DC-coupling case, the detector bias for each channel would have to be added to the background voltage to provide both corrections simultaneously.

Because the focal plane CCDs are cooled, a dynamic range of 80 dB (10,000:1 peak signal/rms noise) is expected. In order to estimate the scene temperature range which can be accommodated by the electronics on the focal plane, we compute the spectral radiance in the 8-12 μm band (from Planck's radiation formula). This indicates that a spectral radiance resolution of $\delta R = 0.022$ watts/m^2 sterrad is equivalent to a 0.05°K temperature resolution at a nominal ambient temperature. At a minimum temperature of $T_{min} = 200°K$; for example, $R_{min} = 4W/m^2$ sterrad. Then for $R_{max} = R_{min} + 10^4 \delta R = 224$, $T_{max} = 475°K$.

Assuming that all gain factors (detector size, IFOV, detector responsivity, input amplifier gain, etc.) are such that R_{max} translates into the maximum CCD well charge, the focal plane electronics is capable of handling an equivalent scene temperature range of 200°-475°K (-100°F to 395°F) at 0.05°K resolution. Raising the minimum temperature to 293°K only raises the maximum temperature to 490°K (68°F-420°F).

This calculation illustrates two important points with regard to background subtraction. First, the dynamic range of cooled CCDs is probably sufficient to permit a constant background level subtraction and still preserve all significant information in a typical scene. Second, the nonlinear relationship between temperature and observable radiant power quickly produces

10

diminishing returns in equivalent temperature dynamic range (at fixed resolution) as the background subtraction level is raised. The off-focal plane electronics outside the dewar will have dynamic range of about 60 dB, and it will lead to a display with a somewhat less dynamic range. Thus, it appears much more important to provide on-focal plane contrast enhancement to ensure that the information available inside the dewar is preserved when it is taken outside.

In summary, a number of PV IR/CCD fabrication and performance issues need to be resolved before answering the question of whether or not AC coupling on the focal plane should be implemented. Examination of these issues by detector and focal plane electronics designers at HRC has begun. In any event, it is probably desirable to incorporate AC coupled electronics off the focal plane to avoid the problems of drift in DC coupled amplifiers. Here, long time constant coupling networks are easily realized, and channel DC levels established on the focal plane can be preserved simply by clamping to zero at the start of every scan line.

Responsivity Equalization

Uniformity of channel-to-channel responsivity is an important requirement if high temperature resolution is to have any meaning. PV detectors can be fabricated at about 5% uniformity, but temperature resolution of order 0.05°K require uniformity of 0.1% or better among channels.

In parallel scanned systems with serial TDI, there is a question about the difficulty of providing a spatially uniform calibration source for the detector array. In this section we examine a signal processing alternative that does not rely upon reference source uniformity.

11

A TDI channel can be treated as a single detector provided that each detector in the channel TDI array views the same spot on the reference (as it does during scanning the scene). Thus, if R_{kj} is the responsivity of the j^{th} detector in the k^{th} channel, B_{kj} is its DC bias, and $P_{kj}(t)$ is the radiant power at time t, the output voltage of the detector is:

$$V_{kj}(t) = R_{kj}P_{kj}(t) + B_{kj}$$

The channel output for (J+1) TDI detectors is

$$V_k(t) = \sum_{j=0}^{J} \{R_{kj}P_{kj}[t - (J-j)\delta] + B_{kj}\}$$

If each detector views the same point in sequence at intervals δ,

$$P_{kj}[t - (J-j)\delta] = P_{kJ}(t) \equiv P_k(t), \ 0 \leq j \leq J$$

Hence, we can write

$$V_k(t) = R_k P_k(t) + B_k$$

where

$$R_k = \sum_{j=0}^{J} R_{kj}; \ B_k = \sum_{j=0}^{J} B_{kj}$$

Channel responsivity equalization can therefore be treated as single detector equalization.

12

The calibration technique assumes that a pair of line sources at different temperatures (T_1, T_2), conveniently located on the field stop for example, are sensed at the end of each line scan. The difference signal,

$$Z_k(t_i) = R_k[P_k(t_i+\delta, T_2) - P_k(t_i, T_1)] + N_k(t_i) \equiv R_k S_k + N_k(t_i)$$

is taken to remove the bias (N_k is the noise). At successive field intervals, τ, the vertical scan mirror is nodded one IFOV so that adjacent detectors have imaged the same reference spot as illustrated in Figure 1.
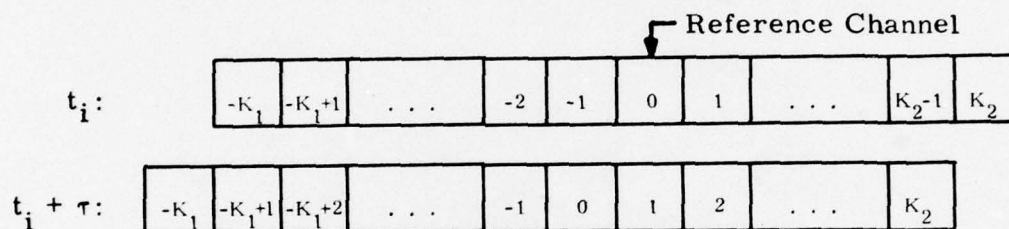
Figure 1. Overscan Geometry

Calibration of the array is performed in pairs--channel 1 relative to channel 0, 2 relative to 1, etc. All channels are referenced to the center channel since errors propagate in direct proportion to distance from the reference channel. A simple iterative algorithm of the form (for $k>0$)

$$R_k(t_{i+1}) = R_k(t_i) + \beta[Z_{k-1}(t_i) - Z_k(t_i+\tau)]$$

$$\beta = c / \sum_{k=0}^{K_2} \overline{Z_k^2}$$

13

permits continuous on-line adaptation to changing environment of both the detectors and calibration source. Although a direct adaptation of the responsivity is indicated, in actuality the gain of the channel output amplifiers would be controlled to ensure equal gains for the full channel electronics.

Figure 2 shows results of a simulation of adaptive gain equalization for 15 elements. The $\log_{10}$ of the error in equalized gain for the first detector (1-dots) and the center detector (8-+ signs) and the last detector (15-solid curve) relative to the reference detector (0) are all plotted for each of 100 iterations (1.7 sec at a field update rate). The S/N is 30:1 (real ratio). The initial detector responsivities for all detectors except 1, 8, and 15 were determined by a random number draw from a Gaussian distribution ($\pm 10\%$ one sigma); and the calibration source was assumed to have 5% one sigma variability from spot-to-spot and look-to-look. Initial gain errors of +30% for detector 1, -30% for detector 8, and +20% for detector 15 were chosen to test the convergence of the algorithm. The iteration parameter constant, c, was decreased linearly from 0.3 down to 0.03 after 50 iterations to speed convergence, thereafter it was held constant at 0.03. After 100 iterations

$$\log_{10} \left| (R_k/R_o) - 1 \right| \leq -3, \ 1 \leq k \leq 15$$

indicating that all channel gains were equalized to within 0.1%.

Experience has shown that this scheme fails for very long arrays. However, by dividing the array into subarrays, the maximum calibration distance to the last detector in any subarray can be minimized. A scheme for equalizing the maximum calibration distance is illustrated in Figure 3 for the case N=10, K=4. The lead detector in each subarray (k_o) is calibrated

14
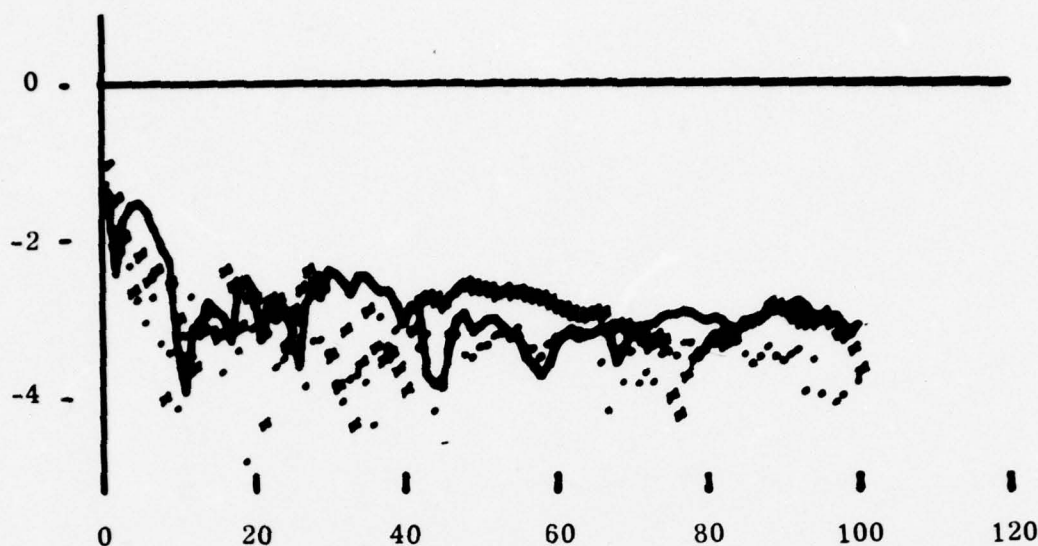
Figure 2. Adaptive Responsivity Equalization
Simulation Results

with respect to the reference detector (0_o), then the other subarray detectors are calibrated with respect to their lead detector. This approach requires K overscans to obtain all the information required for each iteration. In Figure 3, the first overscan provides the single detector shift required for calibration within the subarrays, the remaining provide the overscan of the subarray references; i.e., 3_o with 2_o at $t_i + 2\tau$, 2_o with 1_o at $t_i + 3\tau$, and finally 1_o with 0_o at $t_i + 4\tau$. Although this reduces the iteration rate, different overscans could be performed on successive line scans, rather than successive fields, if less than full parallel scan is implemented. Note that for 100 channels, N=50 and K=10. Thus, no channel would be more than 9 away from the reference, and rapid convergence would be expected based upon the results for K=15 in Figure 2.

15

t_i:				0_0	0_1	0_2	0_3	1_0	1_1	1_2	2_0	2_1	3_0	
$t_i + \tau$:			0_0	0_1	0_2	0_3	1_0	1_1	1_2	2_0	2_1	3_0		
$t_i + 2\tau$:		0_0	0_1	0_2	0_3	1_0	1_1	1_2	2_0	2_1	3_0			
$t_i + 3\tau$:	0_0	0_1	0_2	0_3	1_0	1_1	1_2	2_0	2_1	3_0				
$t_i + 4\tau$:	0_0	0_1	0_2	0_3	1_0	1_1	1_2	2_0	2_1	3_0				

Figure 3.   Equal Calibration Distance Subarrays: $N \leq K(K+1)/2$

As a final comment, PV detectors exhibit good linearity between radiant power and detector current over several decades of incident radiance. This holds true for detectors coupled to CCDs as well, provided that background current bias offset is incorporated to maintain a detector resistance, CCD input amplifier transconductance product greater than one ($R_D g_m > 1$). The dynamic range example discussed previously, indicated a radiance range of only about 55:1 for full 80 dB CCD dynamic range utilization at 0.05°K resolution.   Therefore, it appears that the calibration reference temperature need not be adapted to the ambient-scene temperature to provide a good calibration over the full dynamic range of interest.

PV IR/CCD l/f Noise

It was stated previously that l/f noise in direct-coupled PV IR/CCDs was
negligible. This section presents some recent experimental evidence typical
of hybrid (PbSnTe and HgCdTe) detector/CCDs which supports that state-
ment. Figure 4 is a plot of experimental noise spectral data at $60 \eta A$ CCD
input amplifier drain current. Here again, a high $R_D g_m$ product produces
low drain currents and correspondingly low l/f noise rules. Actually, $30 \eta A$
for $8-12 \mu m$ and $3 \eta A$ for $3-5 \mu m$ are more typical.
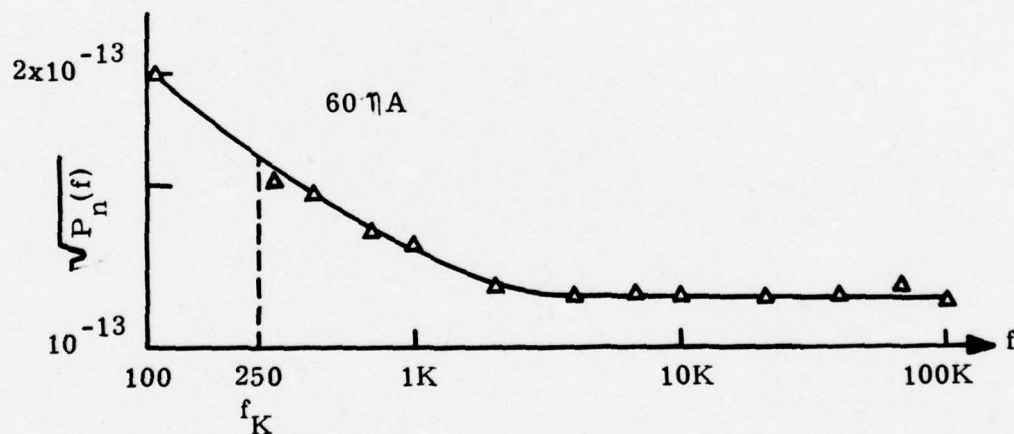
Figure 4. PV Detector/CCD Noise

17

For a noise spectrum of the form

$$P_n(f) = C(f_k/f+1)$$

in a passband of $f_1 \leq f \leq f_2$, the noise variance associated with frequencies below f_k is

$$\sigma_1^2 = C[f_k \ln(f_k/f_1) + (f_k - f_1)]$$

and with frequencies above f_k

$$\sigma_2^2 = C[f_k \ln(f_2/f_k) + (f_2 - f_k)]$$

We know that we cannot take $f_1 = 0$ in this model since $\sigma_1^2 = \infty$ would result. Since $1/f$ noise has been observed down to 10^{-4} Hz (2.8 hour time constants), assume $f_1 = 10^{-4}$ Hz is a lower bound beyond which the low frequency noise essentially has constant spectral density, and hence negligible contribution to the total energy.

In an 875 line TV compatible system with 2:1 interface, the field scan time at 83% scan efficiency is about 13.8 msec. A 100 channel system requires an information bandwidth of about $f_2 = 184$ kHz for 4:3 aspect. At $f_1 = 10^{-4}$ Hz, $f_k = 250$ Hz and $f_2 = 184$ kHz,

$$\sigma_1^2 = (3.9 \text{ kHz})C$$
$$\sigma_2^2 = (185.4 \text{ kHz})C$$

18

Since the total noise standard deviation, $\sqrt{\sigma_1^2 + \sigma_2^2}$, is only about 1% higher than σ_2, it is extremely unlikely that any low frequency flicker noise would be noticeable. In the worst case of a fully parallel scan system, $f_2 \cong 42$ kHz and $\sqrt{\sigma_1^2 + \sigma_2^2}$ is about 4.5% higher than σ_2, indicating $1/f$ noise would still be neglibible at the slowest scan rate.

SECTION II

CONTRAST ENHANCEMENT

INTRODUCTION

This task addresses two related problems in thermal imaging:

1. The dynamic range of the temperatures within a scene (typically 1000:1) can be far higher than the display dynamic range (less than 100:1).

2. Moreover, the scene dynamic range itself can change from scene to scene, requiring the global gain and brightness *controls to be scene* adaptive.

The limited display dynamic range results in the loss of local contrast masking target details. The objective is to reduce the dynamic range over the scene, while maintaining (or even improving) the local contrast in the image. Further, the scene-to-scene brightness range variation should be compensated for by an automatic global gain brightness control. In this section, various local contrast enhancement and automatic gain brightness control techniques being applied to the FLIR imagery are reported.

We assume the display characteristics to be linear between the pedestal brightness B and the dynamic range given by B + R. Given a scene of arbitrary dynamic range R_S and the pedestal brightness B_S, we have to transform the input image to fit the range B and B + R. Since R_S is usually much larger than R, this involves compression of the dynamic range of the

scene. Local contrast would be lost if linear compression were used. Hence, we must make better use of the available grey levels in the display. Second, the relative brightness of distant areas within the scene contributes little to discrimination. Local area brightness controls can help by suppressing slow variations of the brightness across the image, while retaining the local contrast. This concept immediately leads to high-emphasis linear and homomorphic filtering, successfully employed for image enhancement. Indeed, we will establish the correspondence between the two approaches in a later section.

The techniques being investigated are:

- Local area histogram modification

- Local area gain brightness control adaptive to global range variations

- Linear and homomorphic filtering

These techniques are not mutually exclusive. The goal of this study is to define a scheme that combines the best features of each approach applicable in the FLIR context.

HISTOGRAM MODIFICATION

Histogram modification refers to a non-linear mapping of the original intensities to another domain, so that the new intensities have a specified distribution. When the new distribution is uniform, the mapping process is referred to as histogram equalization. The result of histogram equalization is that more frequent grey levels are spread farther so that they occupy

21

a greater portion of the available range, while the less populated grey levels are drawn closer.

The histogram modification routines are based on the following analysis:

Let the original intensity levels be denoted by $x = K_1, \ldots, K_2$. The corresponding histogram normalized into a density function (by dividing by the number of elements in the area), is given by $p_X(x)$, $x = K_1, \ldots, K_2$. We desire a mapping $y = g(x)$ such that y has a given density function $p_Y(y)$, $y = N_1, \ldots, N_2$.

Let the corresponding cumulative distributions be $P_X(x) = \sum\limits_{x = K_1}^{K_2} p_x(x)$

and $P_y(y) = \sum\limits_{y = N_1}^{N_2} p_y(y)$

Since both $P_X(x)$ and $P_Y(y)$ are monotonically increasing, it can be shown that

$$y = g(x) = P_y^{-1}[P_X(x)]$$

Figure 5 illustrates this mapping schematically. A grey level such as x_o is mapped into new intensity y_o by the above transformation. $P_X(x)$ then is the cumulative distribution computed from the desired histogram shape. In Figure 5 we show $P_y(y)$ to be a linear ramp corresponding to a uniform density function. Actually, if the desired density is uniform, it is only necessary to look up the value $P_x(x)$ for a given x. The new intensity $y = g(x)$ is then simply linearly proportional to $P_X(x)$. For an arbitrary

22

distribution $P_Y(y)$ however, the second look-up implied in Figure 5 has to be performed to find the new transformed intensity y_o.   Figure 5 also shows the way for an efficient look-up table implementation of histogram modification using CCDs.
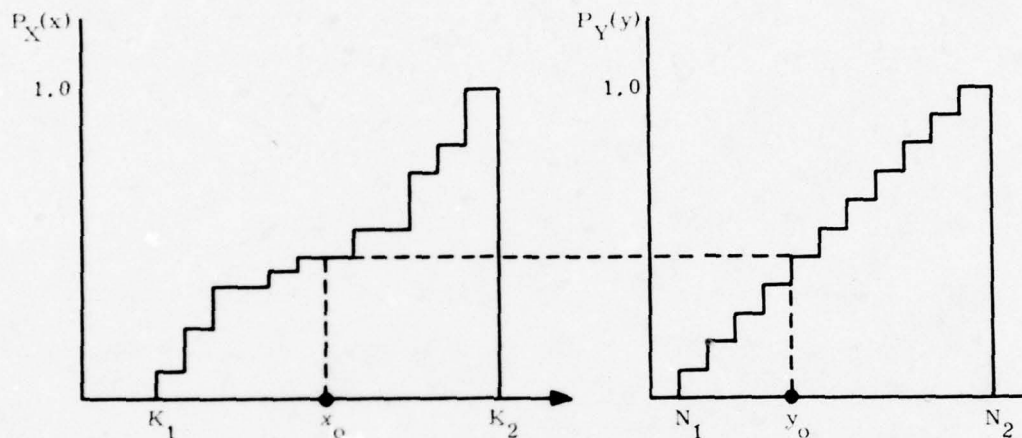
Figure 5. Histogram Modification Mapping

Full Frame Histogram Modification

Full frame histogram modifications with uniform (equalization) and hyperbolic distributions have been tried. Figure 6 is an 8-bit (256 levels) quantized NVL-supplied image reproduced on an Optronics film writer constrained to a 50:1 dynamic range. Figure 7 is the corresponding histogram. Note that there are only 163 levels in the range. This image was then histogram equalized so that the resultant histogram would be uniform in the range 0-255. Figure 8 is the corresponding equalized image.

23
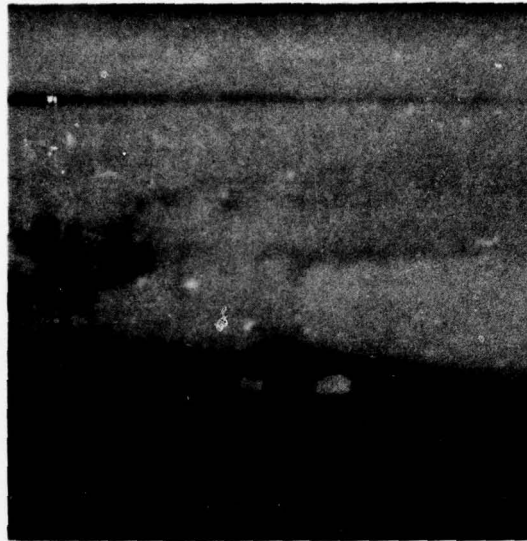
Fi   6. Original Image with 163 Grey Levels
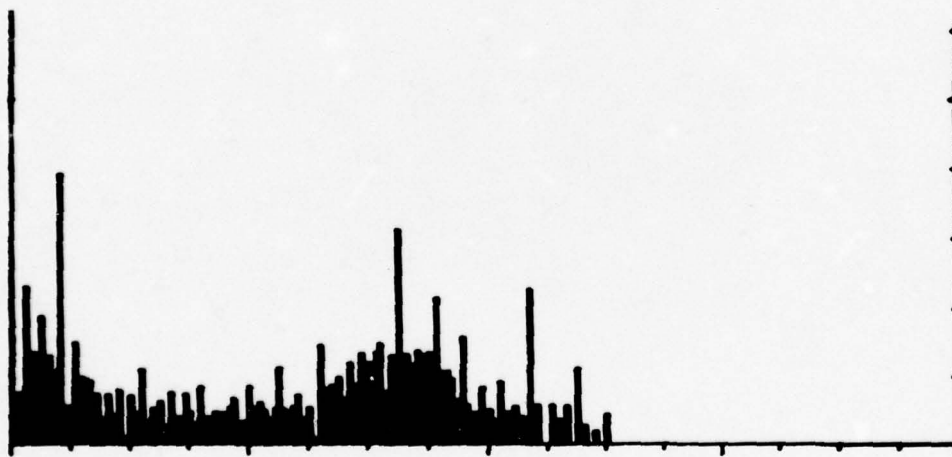
Figure 7. Histogram of Image in Figure 6

24
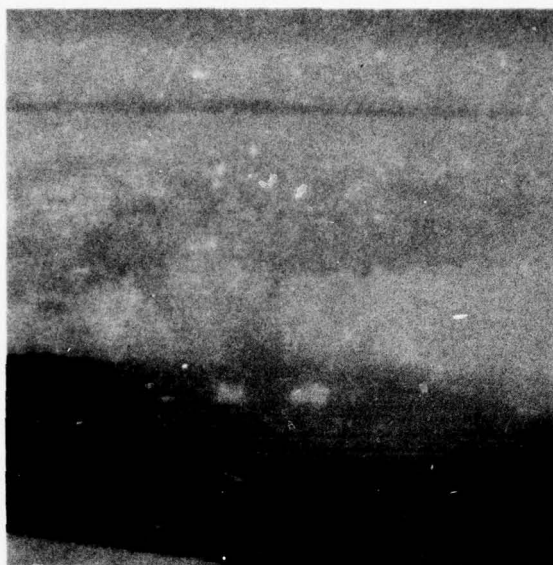
Figure 8. Full Frame Histogram Equalized Image

An alternative to equalization is histogram hyperbolization (reference 1). This is based on the premise that the perceived brightness b is proportional to the logarithm of the display intensity I. Therefore, the observed brightness levels (log I) should be uniformly distributed. To achieve this, the distribution of the displayed intensities should be hyperbolic, instead of uniform. Figure 9 is a reproduction of the image modified in this manner. Histogram hyperbolization appears to be an improvement over equalization.

Local Area Histogram Modification

In local area histogram modification the image points in a small area (window) are modified using a transformation computed from the histogram of the larger subregion surrounding the local area. Figure 10 illustrates the inner and outer regions. Local area histogram equalization might alle-

25

Figure 9. Full Frame Histogram Hyperbolized Image

Figure 10. Illustrates Inner and Outer Windows for
Local Area Histogram Modification

26

viate some problems of full frame equalization, because the grey scale transformation is now based on the area immediate to the area being equalized. However, the following questions arise in local area histogram equalization:

The limits of the equalized histogram (N_1, N_2) should be determined for each local area. If all local area histograms ar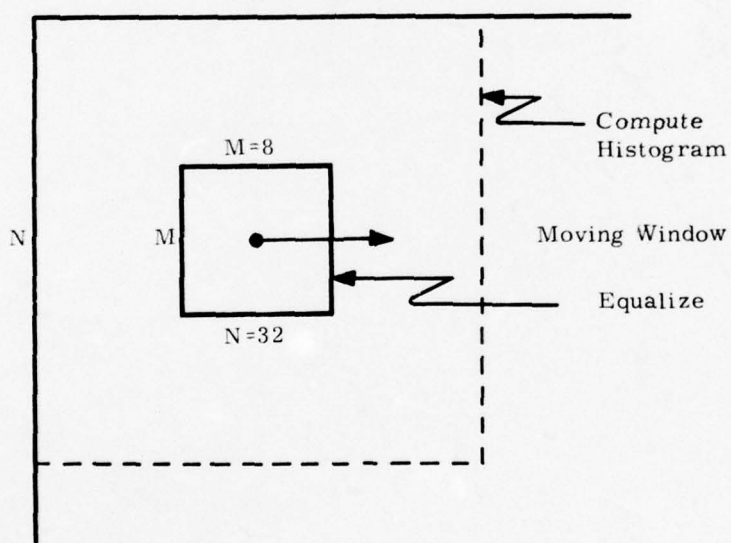e equalized out to the same intensity limits as in Figure 11, the relative contrast between two areas with totally different brightness is completely lost and undesirable texture is brought out in bold relief as shown in Figure 12. Indeed, this approach corresponds to completely removing the low spatial frequency components in the image--clearly an unwarranted phenomenon.

An alternate approach is to make the new limits (N_1, N_2) directly proportional (Figure 11b) to the old limits (K_1, K_2) in each local area as follows:

$$N_1 = \alpha K_1 + \beta \text{ and } N_2 = \alpha K_2 + \beta$$

where α and β are chosen so that the maximum and minimum transformed intensities (over all local areas) correspond to the extreme limits of the display - (B, B + R). Figure 13 was equalized in this manner, and we see that the relative contrasts between areas of different brightness are preserved. However, note the boxy anomaly around the hot target, where the severe banding effect has been accentuated. This is because the local area histogram is strongly bimodal on a target/background boundary as in Figure 14. The higher intensity peak corresponds to the hotter target and the lower, broader peak to the background. The transformed intensities are shown below the histogram. We note that the equalization compresses the valley region between the two peaks in the histogram. This raises the

27
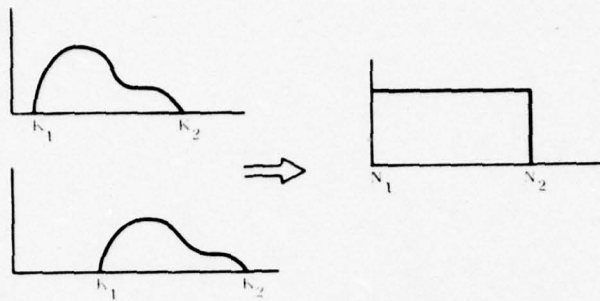
a. Same Interval (N_1, N_2) for all Local Areas
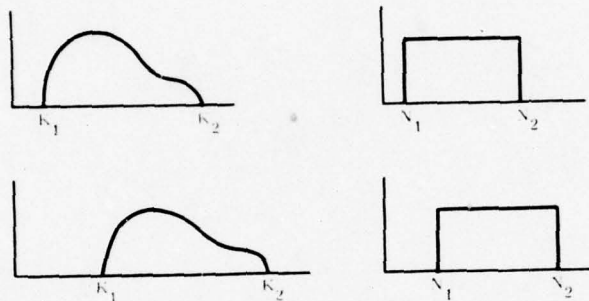
b. Proportional Intervals
$$N_1 = \alpha K_1 + \beta, \qquad N_2 = \alpha K_2 + \beta$$
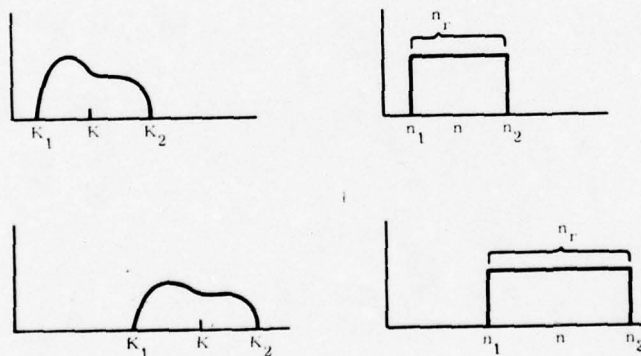
c. $n_r \, \alpha \, n$ $n = A(K - K_g) + Bias$

LAGBC Incorporated into the Equalization Intervals

Figure 11. Different Equalization Interval Schemes

28
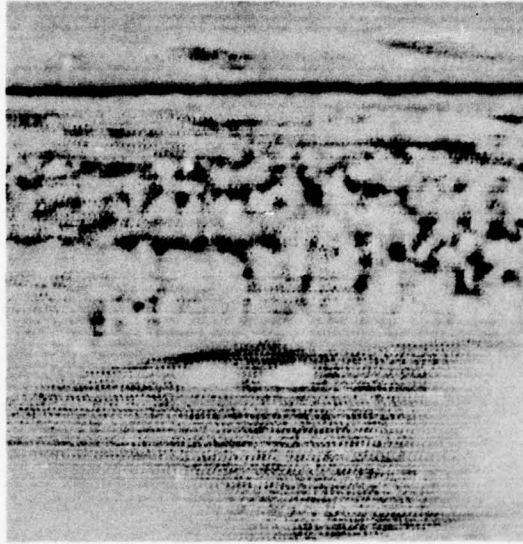
Figure 12. Local Area Histogram Equalized, as in Figure 11a

Figure 13. Local Area Histogram Equalized, as in Figure 11b
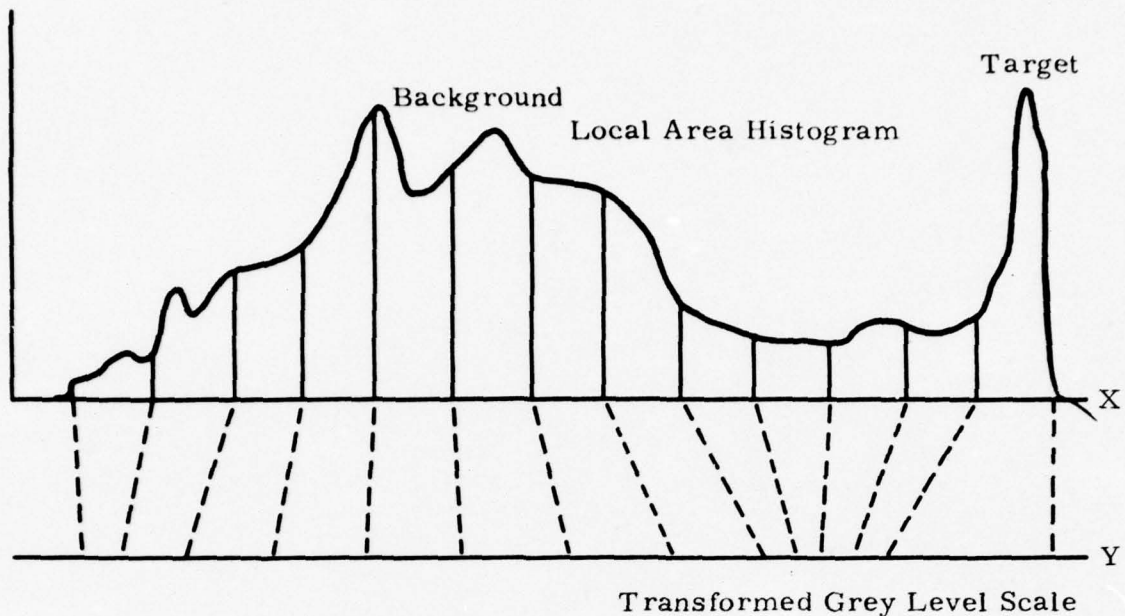
Figure 14. Shows Bimodal Distribution on a Local
Area with Target and Background

intensity level of middle intensity background region (where banding has
occurred) closer to the hot target, accenting the banding effect. To alleviate
this, other shapes, including the bimodal catenary, are being investigated
for histogram modification. However, the gain for each local area $\dfrac{N_1 - N_2}{K_1 - K_2}$
is not spatial frequency dependent as it was with Figure 12. We
would still like to attenuate low spatial frequencies in the image (without
removing them altogether). This would make more room in the display
dynamic range for local contrast to be enhanced.

A third approach is to make the ranges (N_1, N_2) adaptive over each local
area in a manner precisely analogous to the local area gain brightness
approach to be outlined in the next subsection. We will defer the discussion
of this equalization until then.

30

LOCAL AREA GAIN BRIGHTNESS CONTROL

Given a scene that has a dynamic range between B_S and $B_S + R_S$, local area gain brightness control should:

- Vary local average brightness to contract the overall dynamic range of scene

- Enhance local variations above the human contrast sensitivity threshold to bring out the details

- Automatically fit the new extremes in the image to the display limits B and $B + R$

The above assumes that the dynamic range in a scene is larger than that of the display. It is of course possible to have a low dynamic range image where it is not necessary to compress the overall dynamic range to fit the display. Therefore, this feature should be scene adaptive. Following is a formulation that does precisely this:

The image intensity I_{ij} of each point is transformed based on local area statistics--the local mean m_{ij} and standard deviation σ_{ij} computed on a local area surrounding the point. The new intensity,

$$\hat{I}_{ij} = \underbrace{G_{ij}(I_{ij} - m_{ij})}_{(1)} + \underbrace{A(m_{ij} - \overline{m}) + H + B}_{(2)}$$

and the local gain,

$$G_{ij} = \frac{C_s}{\alpha \sigma_{ij}} \; [A(m_{ij} - \overline{m}) + H + B]$$

31

M is the global mean and H is a bias value computed to restore the average value to the image. A is constant for a given scene. A and H are determined by the actual dynamic range present in the image.

This equation, based upon the model shown in Figure 15a, shows that the transformed image is composed of a slowly varying brightness component (2) and superimposed higher frequency local variations (1). The local gain operates on the first term and is itself proportional to the average local display brightness component (2). As seen in Figure 15b, the gain A reduces the slowly varying component bringing it closer to the global mean, so that there is more room for enhancing local details.

The rationale for making the local gains G_{ij} proportional to the local display brightness is that over a large range of intensities, the minimum resolvable intensity difference ΔI is directly proportional to the average local intensity I. The contrast sensitivity C_s is the proportionality constant. Also the local gain is made inversely proportional to the local standard-deviation effort to enhance areas of little contrast.

Note that the above formulation is similar to that of reference 2, but with the following constraints.

- The gain A and the bias H are computed on a per scene basis to fit the modified scene range to the display. A can therefore be greater than unity (for scenes with low dynamic range)

- The local gains are now directly related to the contrast sensitivity of the human eye

- Standard deviation determines the extent to which the local standard deviation governs the local gain

32
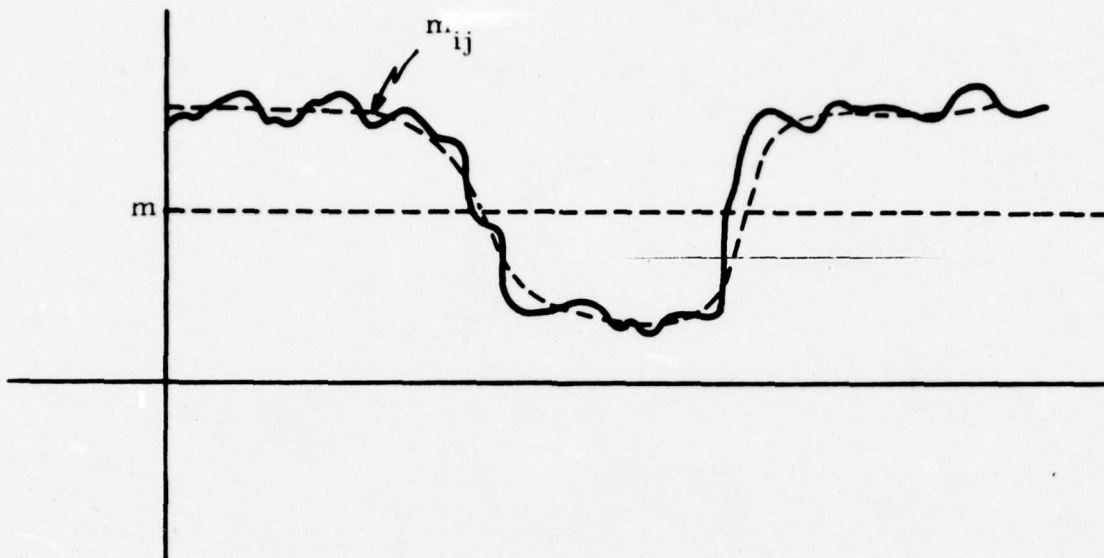
Figure 15a. Original One-Dimensional Image
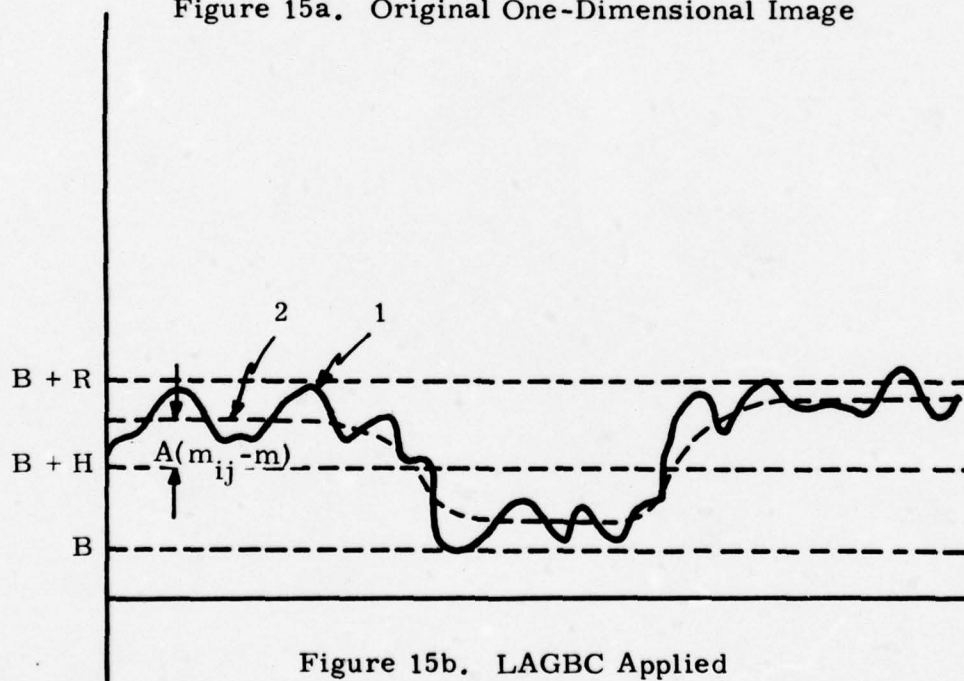
Figure 15b. LAGBC Applied

1. High Frequency Component
2. Low Frequency Component

The size of the local area over which the local means are determined is important. This determines the spatial frequencies that are attenuated during the LAGBC process. An exact correspondence between this and spatial frequence filtering will be established in a later section.

Implementation

The primary computation involved is in evaluating the local means and standard deviation over a sliding window centered at each point. This is inefficient from the implementation viewpoint. Two alternate approaches are being implemented:

1. Similar to the overlapping window approach used in Local Area Histogram Equalization (Figure 10). The mean and standard deviation computed for the outer window are assumed constant over the inner window. We call this the piecewise constant approach.

2. The second approach is to have nonoverlapping local regions. The local mean and standard deviation at each point are obtained by a simple two-dimensional interpolation. This is a piecewise linear approximation.

Local Area Gain and Brightness Applied to Histogram Equalization

Figure 11c illustrates this approach where the limits of equalization for each area histogram are determined by considerations similar to that of LAGBC.

By comparison with the LAGBC equation, we see that,

- Mean local display intensity $\simeq \bar{n}$
- Local gain is approximated by $\dfrac{n_1 - n_2}{K_1 - K_2}$ factor

The local gain should be proportional to the local mean brightness; i.e.,

$$\frac{n_1 - n_2}{K_1 - K_2} = \alpha_1 \frac{\bar{n}}{\sigma_{Local}} C_s$$

If we approximate σ_{local} by $\alpha_2(K_1 - K_2)$, the spread of the LAH, we get

$$n_r = n_1 - n_2 = C_1 \bar{n}$$

which implies that the range of the mapped intensities should be directly proportional to the average mapped intensity for that local area. Also, by comparison with the LAGBC equation, the average local display brightness is

$$\bar{n} = A(\overline{K} - \overline{K}g) + H + B$$

where $\overline{K}g$ is the global intensity mean and B is the pedestal brightness as before. A is the brightness gain and H is a bias, computed as before, so that the global extrema of the mapped intensities correspond to the limits of the display (B, B+R). This approach is being implemented and preliminary experimentation with different values of the proportionality constant C_1 is being made.

LINEAR AND HOMOMORPHIC FILTERING

High frequency emphasis filtering of the FLIR imagery has a two-fold purpose:

- Attenuating the low spatial frequencies reduces the global dynamic range, without sacrificing local contrast

- High frequency emphasis crispens edges and other fine detail that might not have been otherwise visible.

The above concept is also evident in the local area gain brightness model, where the local brightness envelope (see Figure 15a) denoted by the broken line, comprises the low frequency components upon which the higher frequency local variation is superimposed. Attenuating the low frequencies compresses the low frequency envelope thereby reducing the overall dynamic range. Of course, removal of the low frequency components is not desired. The corresponding filter structure and the overall frequency response is shown in Figure 16a and b. The low frequency gain A_ℓ, the high frequency gain A_h, and the 3dB frequency f_{3dB} of the low pass filter should be tuned to the imagery.

Homomorphic filtering is implied in the LOG and EXP boxes. Homomorphic processing assumes a multiplicative model of the image (reference 3), which is true of imagery in the visible and near IR parts of the spectrum where the image can be modeled as a product of low frequency illumination and higher frequency reflective components. The multiplicative model does not seem to be directly applicable for FLIR imagery, but can be easily incorporated in the present filter structure.

36
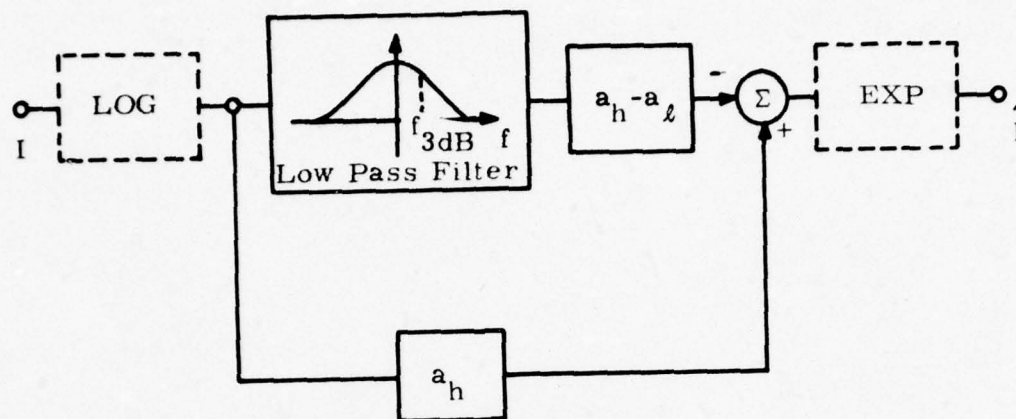
Figure 16a.   Basic High Pass Filter With
Optional Homomorphic Processing
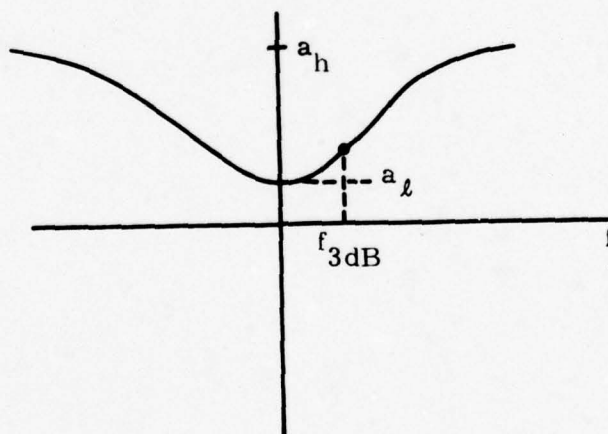
Figure 16b.   Overall Frequency Response of
Filter in 16a

37

The Low Pass Filter

The basic low pass filter used in Figure 16a can be realized in different ways. We are designing two basic realizations: (1) direct convolution finite impulse response (FIR) Gaussian filters and (2) recursive infinite impulse response (IIR) filters.

Gaussian Filter (FIR)

The low pass filter has the following transfer function:

$$H(f_x, f_y) = \text{Exp} \left[- \frac{f_x^2 + f_y^2}{2\sigma^2} \right]$$

and the corresponding FIR convolution filter also has a Gaussian shape:

$$h(x, y) = \frac{1}{2\pi\sigma'^2} \text{Exp} \left[- \frac{x^2 + y^2}{2\sigma'^2} \right]$$

where

$$\sigma' = \frac{1}{2\pi\sigma}$$

This filter does not have a sharp frequency transition but has the advantage that the spatial impulse response can be truncated to a finite area, without Gibb's phenomenon occuring. Also, the filter is isotropic; i.e., it possesses circular symmetry. From the realization standpoint, it has the advantage of being separable; i.e.,

$$h(x, y) = h_x(x) \, h_y(y)$$

38

Therefore, a two-dimensional convolution can be cast as two sets of successive one-dimensional convolutions:

$$y(m, n) = \sum_{k=0}^{M-1} \sum_{l=0}^{M-1} x(k, l) \ h(m-k, n-l)$$

$$= \sum_{k=0}^{M-1} h_x(m-k) \sum_{l=0}^{M-1} x(k, l) \ h_y(n-l) \quad \text{where } m, n = 0, \ldots, N$$

A direct two-dimensional convolution of a MxM filter with a NxN image would require $N^2 M^2$ multiplications and additions. This separability enables us to filter with one-dimensional convolution of the rows of the image followed by one-dimensional convolution of the resultant columns as implied above. This requires only $2N^2 M$ multiplications and additions--a distinct advantage over the $N^2 M^2$ requirement, especially for large filters. However, the filtered rows have to be transposed, so that the columns can be filtered--a disadvantage when working with a sequential mass medium such as magnetic tape. Therefore, we are implementing both the direct 2-D spatial filter for small filter sizes, and the separable filter for wide impulse response filter realization.

Choice of 3dB Frequency

The choice of 3dB frequency of the basic low pass filter in Figure 16a is very important in contrast enhancement. If only edge emphasis/crispening is desired, the corresponding filter can only be a few pixels wide. The corresponding frequency response is shown in Figure 17a where the 3dB frequency is very high. On the other hand, for attenuating the low frequency components for dynamic range compression, we should have a much lower
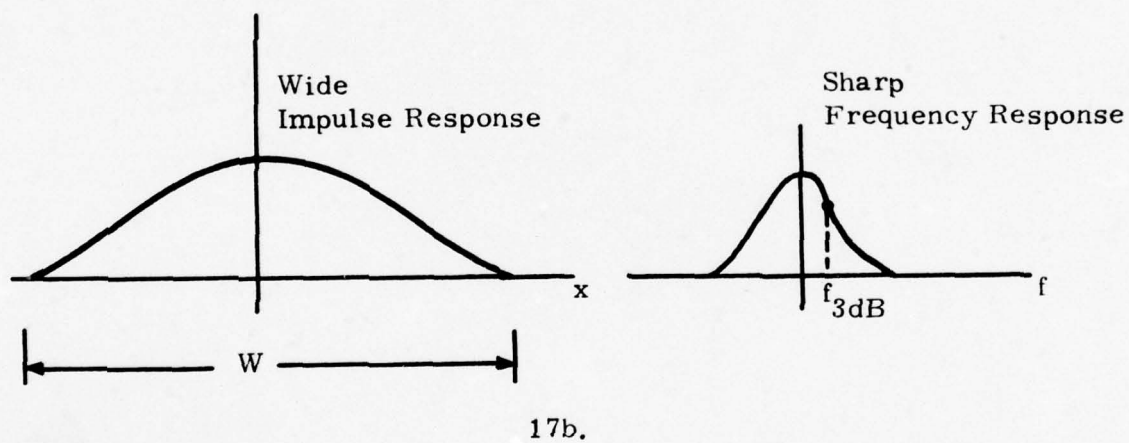
39
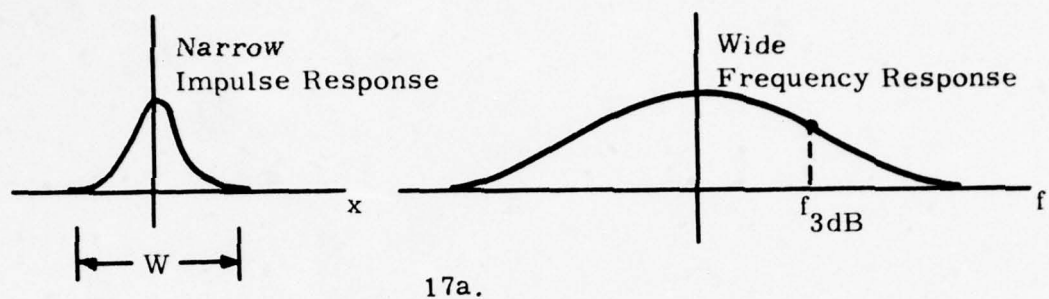
Figure 17a and b. Illustrating the Relationship between Filter Size and 3dB Cutoff Frequency for a FIR Filter

3dB cutoff in the filter, Figure 17b, which requires much larger spatial filters. Unless the fast Fourier convolution is used, direct convolutional filtering is computationally expensive for large filters. Hence, the separable Gaussian filter will be useful for this implementation.

Recursive Filters

The large impulse responses needed for spatial filtering for contrast enhancement can be realized using simple two-dimensional recursive filters. The advantages of recursive filters are:

- Computational efficiency--any desired 3dB frequency (corresponding to large impulse responses) is realizable with a low order recursive filter.

- A two-dimensional recursive filter can be designed as a product of two one-dimensional low pass filters. The resultant filter, although not isotropic, may be quite adequate for contrast enhancement (reference 4).

- Ease of implementation--looking ahead to 2nd Generation FLIR implementation, low order (1 or 2) two-dimensional recursive filters are readily implementable using present CCD technology. Only 2n line buffers are needed for an n^{th} order recursive filter. It is expected that n=1 or 2 will be sufficient for the present application.

Low order, two-dimensional recursive Butterworth filters are being programmed to this end. If these filters prove to be of value for contrast enhancement, they will provide ready CCD implementation.

Relation Between LAGBC and Linear Filtering for Contrast Enhancement

The basic local area brightness control shifts the local means closer to the global mean to reduce the overall dynamic range; the contrast term enhances the local contrast. This was illustrated in Figure 15. As pointed out before, this corresponds to attenuating the low spatial frequencies and enhancing the high spatial frequencies. This can be seen as follows. We recognize that the local mean is the convolution of the rectangular window with the image. In the spatial frequency domain, this corresponds to filtering with a low pass filter that is a product of two Sinc functions; i.e.,

$$H(f_x, f_y) = Sinc(f_x) \; Sinc(f_y)$$

We may as well replace this filter by a Gaussian (or other) low pass filter with the same noise equivalent bandwidth. Figure 18 is a realization of the LAGBC using such a low pass filter instead of the averaging process. The frequency response at each stage of the process is shown in the insets. For dynamic range compression, the gain A is less than unity and the high frequency gain G is greater than unity. The resultant response is much like the basic high frequency emphasis filter described above. The spike at the origin corresponds to the bias added to restore the average brightness. There are differences; for example, the high frequency gain G is a function of the low frequency component which makes the process non-linear. Nevertheless, the remarkable similarity between the two approaches leads us to believe that they are equivalent in function. This exercise serves a two-fold function:

- Relates the heuristic LAGBC concept to the rigorous linear filter theory. The cutoff frequency of the low pass filter--the key parameter--determines the size of the averaging window

42

in LAGBC and vice versa. The low frequency gain A is
determined by the range of the scene, and the high frequency
gain G, by the average local intensity.

● Ease of implementation--computation of local mean in LAGBC
(especially over large windows) may be computationally unfea-
sible. In the equivalent LPF realization, Figure 18, we can
implement a simple two-dimensional first order recursive
filter which can be readily implemented with present CCD
technology.
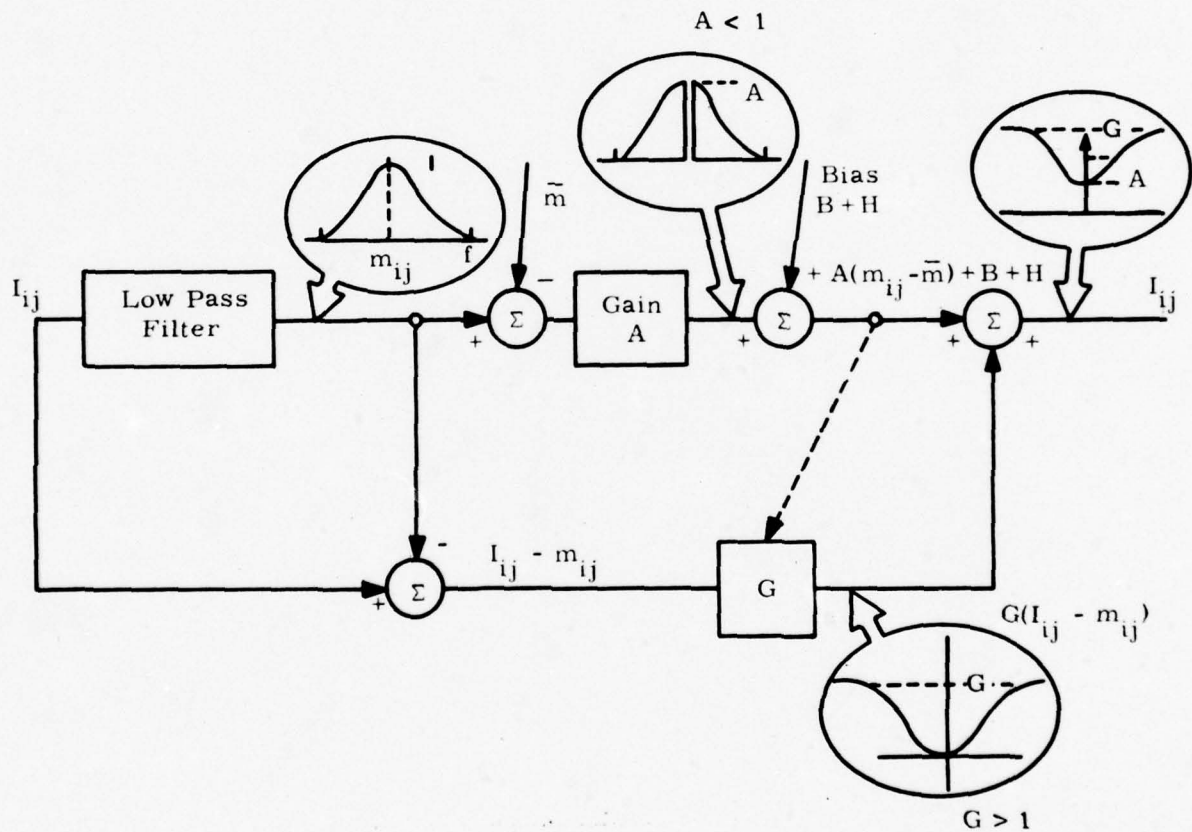
Figure 18. A Linear Filter Approach
to Gain-Brightness Control

43

REFERENCES

1. W. Frei, "Image Enhancement by Histogram Hyperbolization," Unpublished manuscript, Department of Electrical Engineering, University of Southern California.

2. D. Ketcham, et al, "Image Enhancement Techniques for Cockpit Displays," Hughes Report No. P74-530R.

3. T. G. Stockham, "Image Processing in the Context of a Visual Model," Proc. IEEE, Volume 60, pp. 828-843, July 1972.

4. E. L. Hall, "A Comparison of Computations for Spatial Frequency Filtering," Proc. IEEE, Volume 60, pp. 887-891, July 1972.

SECTION III

MRT ENHANCEMENT

There are two approaches to improving system sensitivity by image pro-
cessing: intra-frame and inter-frame spatial averaging. In intra-frame,
smoothing one trades resolution for enhanced S/N. In inter-frame
smoothing the tradeoff is basically between data rate and S/N, although
resolution will be lost if registration is not maintained between stacked
image frames. To date, only intra-frame smoothing has been addressed
in the program. Median, hysteresis and variable width spatial filtering are
currently being investigated.

The median filter simply computes the median intensity in a small sliding
window (3x3 or 5x5) and replaces the pixel intensity at the center with the
median value for the window. This filter has the advantage that median
values are measured values, hence original data is almost preserved,
albeit rearranged. It appears to be most useful as a pre-filter operation
to remove large noise spikes.

Hysteresis filtering is a method which preserves significant extrema in an
image while smoothing out local fluctuations. Two-dimensional hysteresis
filtering is achieved simply by applying one-dimensional filters sequentially
in orthogonal directions (separability assumption). For one dimension, the
algorithm is:

45

If $(k=1)$, $y_k = x_k$

If $(y_k < x_{k+1} - L/2)$, $y_{k+1} = x_{k+1} - L/2$

If $(y_k > x_{k+1} + L/2)$, $y_{k+1} = x_{k+1} + L/2$

Otherwise, $y_{k+1} = y_k$

Here, x_k is the original pixel intensity, y_k is the filtered intensity, and L is the hysteresis window width.

A number of variable width spatial filter algorithms have been proposed, differing mainly in the criteria used to select the filter size. Examples are criteria based upon the local gradient; or the texture, to separate object boundaries and avoid smoothing across them; or upon the local average intensity to make the smoothing window S/N dependent for quantum noise limited pictures.

In FLIR imagery processing, object edge detection is the most attractive criteria because object shape and contrast appear much more useful than texture as classification clues. Also, although FLIR detectors are generally BLIP, the actual range of S/N in the scene does not vary markedly. Hence, local S/N variation is not a good criteria for chosing local filter width. The range of filter width selection should depend upon the average S/N in the scene, however, if effective smoothing is to be accomplished.

Use of the gradient as a criterion for selection of a variable filter width to avoid smoothing across edges leaves something to be desired. The reason is that no error in smoothing an edge results if the edge gradient is constant within the width of the smoothing window.

46

To investigate this further, assume that an edge has the form of a plane ridge, at least over the smoothing window area. If the x-coordinate is in the direction of the gradient, the edge can be described by

$$f(x, y) = a + bx + cx^2 + dx^3$$

The smoothed intensity (for a box filter) is then

$$F(x, y) = \frac{1}{L^2} \int_{x-L/2}^{x+L/2} du \int_{y-L/2}^{y+L/2} dv \, f(u, v) = (a + \frac{cL^2}{12}) + (b + \frac{dL^2}{4})x + cx^2 + dx^3$$

For specified error in the smoothed value at the center of the window

$$\left| \frac{F(0, 0) - f(0, 0)}{f(0, 0)} \right| \le \epsilon \quad \text{if} \quad \frac{L^2}{12} \left| \frac{c}{a} \right| \le \epsilon$$

Thus, the smoothing error does not depend upon the local gradient, b, but rather on the local curvature, c. Furthermore, we now have a criterion for selecting the smoothing window size to achieve specified error.

This analysis suggests the following approach. At each pixel, the least square error estimate of the coefficients a and c in $f(x, y)$ is computed under the assumption that x is in one of the four principle directions (0°, 90°, $\pm 45^{\circ}$). For a 7 pixel, 1-D window these are

$$a = (7g_a - g_c)/21 \qquad c = (g_c - 4g_a)/84$$

where

$$g_a = \sum_{i=-3}^{3} f_i \qquad g_c = \sum_{i=1}^{3} i^2 (f_i + f_{-i})$$

$$f_i = f(x+i\Delta x, y)$$

47

Use of a 7 pixel window provides the advantage of some noise smoothing in the filter width selection criterion.

Having computed a and c for each of the 4 principle directions, the maximum L for which $|c|L^2 \le 12\epsilon|a|$ is found in each direction. Then the minimum of the max L for the filter width is chosen. A variation on this scheme is to orient an asymmetric filter along the edge by allowing the filter width in the direction normal to that of the min max L to remain the max L for that direction.

These variable width filters are being investigated using 2-D Gaussian filters of equivalent rectangular width L=0, 2, 3, 5, and 7. These correspond to S/N gains against white noise of 0, 3, 5, 7 and 8.5 dB, respectively.

Figure 19 shows results of the addition of median and hysteresis filtering of a FLIR image with white Gaussian noise at an average S/N = 5dB (peak signal to rms noise). The dark lines are the result of a software anomaly which caused the Optronics film writer occasionally to skip lines. In Figure 19, A is the original image, B is the noisy image, C is the median filtered image (5 x 5 window) and D is the hysteresis filtered image (L = 8). Both filters preserve target edges, but do not reduce the grainy quality of the picture appreciably. Computation of the mean square error between the original and the noisy, median and hysteresis filtered images, respectively, yields an overall S/N gain of about 4 dB for the median filter, and less than 1 dB for the hysteresis filter. These measures correlate with the differences in subjective quality of the two filtered images.
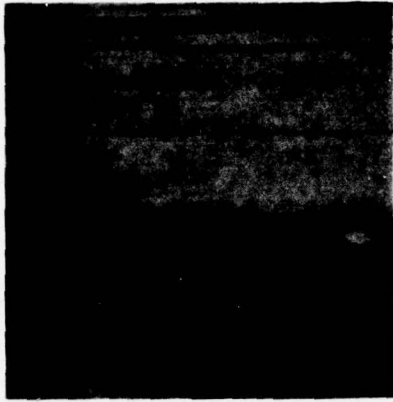
A

B
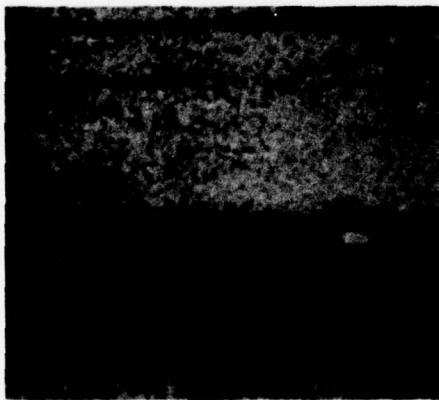
C

D

Figure 19. Median and Hysteresis Filtered FLIR Images

49

SECTION IV

RESOLUTION RESTORATION

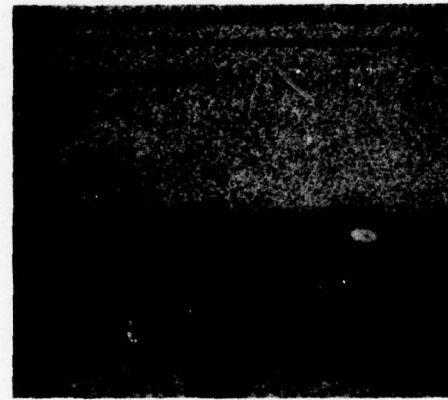Work on full frame focus restoration and super resolution is in its preliminary stages. Space variant blur functions have been modeled and images blurred for subsequent restorations. However, no 2-D restoration algorithms are operational at this time.

The question of depth of field has a bearing on the OTF model for resolution restoration studies. In particular, if the autofocus subsystem can achieve focus at the hyperfocal distance, depth of field problems are greatly alleviated.

Exact geometrical optics analysis gives the maximum range, r_1, and minimum range, r_2, for which the diameter of the angular blur circle is w as:

$$r_1 = \frac{r}{1-(wr/D)} \quad \text{and} \quad r_2 = \frac{r}{1+(wr/D)}$$

where r is the range at which the system is in focus and D is the aperture diameter. At the hyperfocal distance

$$r_o = D/w$$

the depth of field extends over $r_o/2 \leq r \leq \infty$. For example, with a NFOV 8-inch aperture and w=0.05 mrad, the depth of field extends from 2km out to infinity. With a WFOV 2-inch aperture and w=0.2 mrad, $r_o/2=127m$. This implies that only the relatively slight space variant off-axis blur of the OTF need be restored if system focus at the hyperfocal distance is maintained.

50

In case $r \neq r_o$, the depth of field δ is

$$\delta = r_1 - r_2 = \frac{2wr^2/D}{1-(wr/D)^2} \xrightarrow[r<<D/w]{} \frac{2wr^2}{D} = \frac{2w^2 r^2 F^\#}{a}$$

for a detector of width "a" and IFOV $\approx w$. Thus it is only at very short ranges where the depth of field is severely constrained by F# and resolution.

SECTION V

STATISTICAL ANALYSIS OF IMAGERY

Statistical analysis of FLIR imagery is being performed to better quantify the characteristics of thermal imagery, and to provide quantifiable measures of the quality of enhancement processes. The statistics being extracted can be categorized as shape, intensity and texture features. Each is discussed separately below.

SHAPE STATISTICS

All shape statistics are derived from an object-boundary tracing routine and the resulting chain code which describes the perimeter of the object. The boundary extraction routine uses the Sobel gradient operator. Referring to Figure 20, the Sobel gradient operator is:

$$G_s(E) = S_x + S_y$$

where

$$S_x = |AA-AB| \quad S_y = |AL-AU|$$

$$AA = C + 2F + I, \quad AB = A + 2D + G$$
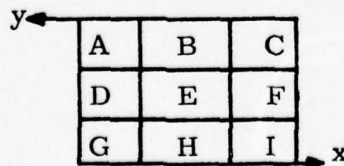
$$AL = G + 2H + I, \quad AU = A + 2B + C$$

Figure 20. Sobel Gradient Geometry

52

To find a first point on the boundary, an initial point P_0 (Figure 21) on the outside of the target is specified, along with the direction D_0, to proceed so as to intercept the boundary. At a point Q_0 the Sobel gradient exceeds a predetermined threshold. Rather than start tracking the boundary at this point, however, two new paths, P_1 and P_2, parallel to and stradling P_0, are chosen and points Q_1 and Q_2 are found on the boundary similar to Q_0.
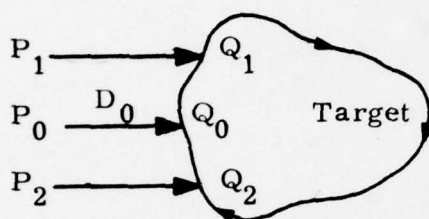
Figure 21. Initial Boundary Search

From Q_0, Q_1 and Q_2 and their eight neighbors, we compute an average boundary intensity and an average boundary Sobel gradient. Q_0 is chosen as the starting point for tracing the boundary in a clockwise direction. The next boundary point is determined from the values of G, Sx, Sy, AA, AB, AL and AU of that point, as well as the average boundary intensity and Sobel gradient following the logic illustrated in Figure 22.

This procedure does not necessarily terminate at the starting point Q_0. Typically the boundary closes as below. By discarding the leading edge $\overline{Q_0 R_0}$, a
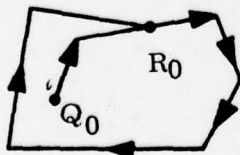
closed boundary starting and ending at R_0 is generated.

53
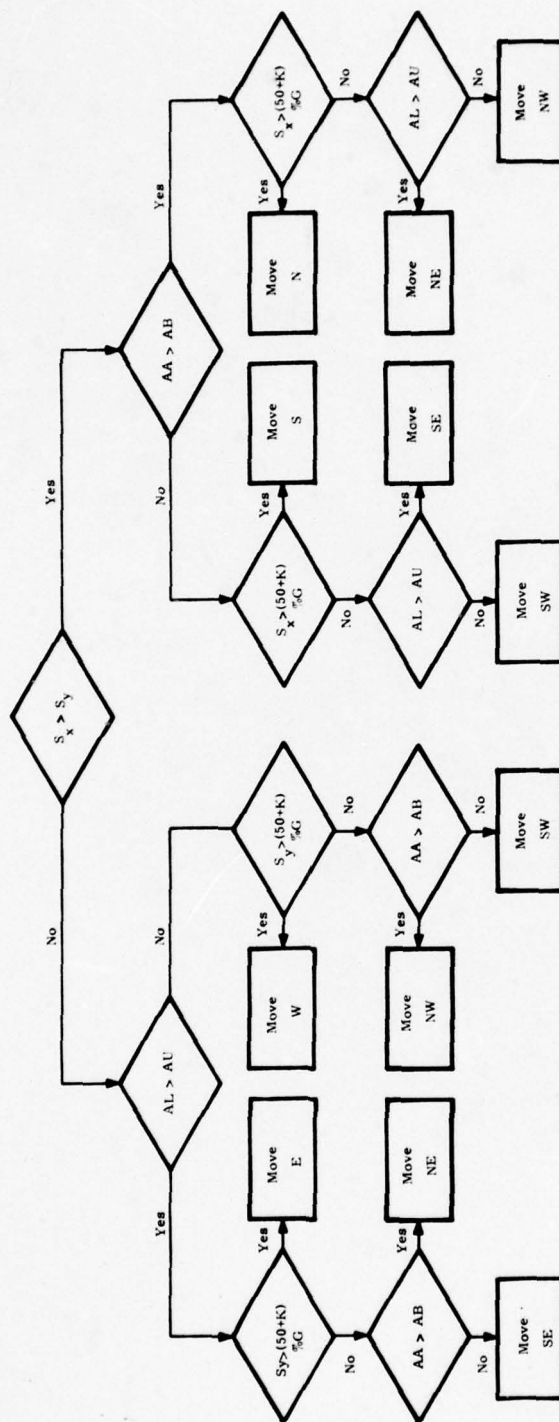
Figure 22.   Boundary Tracing Logic

54

Figure 23 shows a line printer plot of a tank extracted by this technique. Pixel intensities interior to the target are negated for ease of separating target from background when extracting brightness and texture features. The notation (10x20/137) gives the (Height x Width/Area) of the target in pixels.

Figure 24 shows an example of the difference in the extracted boundary between unfiltered and a median filtered image. Median filtering helps to eliminate edge anomalies which can on occasion cause the boundary detector to trace the edge one pixel into the interior.
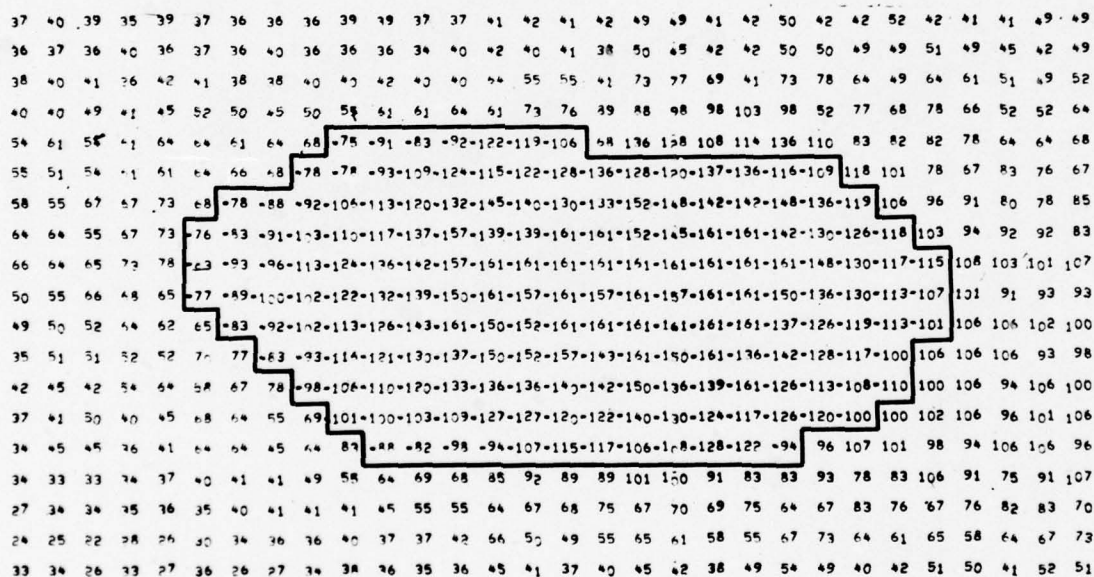
```
37  40  39  35  39  37  36  36  36  39  39  37  37  41  42  41  42  49  49  41  42  50  42  42  52  42  41  41  49  49
36  37  36  40  36  37  36  40  36  36  36  34  40  42  40  41  38  50  45  42  42  50  50  49  49  51  49  45  42  49
38  40  41  36  42  41  38  38  40  43  42  40  40  54  55  55  41  73  77  69  41  73  78  64  49  64  61  51  49  52
40  40  49  41  45  52  50  45  50  55  61  61  64  61  73  76  39  88  98  98 103  98  52  77  68  78  66  52  52  64
54  61  55  41  64  64  61  64  68 -75 -91 -83 -92-122-119-106  68 136 138 108 114 136 110  83  82  82  78  64  64  68
55  51  54  41  61  64  66  48 -78 -78 -93-109-124-115-122-128-136-128-120-137-136-116-109 118 101  78  67  83  76  67
58  55  67  67  73  68 -78 -88 -92-106-113-120-132-145-140-130-133-152-148-142-142-148-136-119 106  96  91  80  78  85
64  64  55  67  73 -76 -43 -91-103-110-117-137-157-139-139-161-161-152-145-161-161-142-130-126-118 103  94  92  92  83
66  64  65  73  78 -43 -93 -96-113-124-136-142-157-161-161-161-161-151-161-161-161-161-148-130-117-115 108 103 101 107
50  55  66  48  65 -77 -89-120-102-122-132-139-150-161-157-161-157-161-161-150-136-130-113-107 101  91  93  93
49  50  52  44  62  65 -83 -92-102-113-126-143-161-150-152-161-161-161-141-161-161-137-126-119-113-101 106 106 102 100
35  51  51  52  52  74  77 -63 -93-114-121-130-137-150-152-157-143-161-150-161-136-142-128-117-100 106 106 106  93  98
42  45  42  54  64  58  67  78 -98-106-110-120-133-136-136-140-142-150-136-139-161-126-113-108-110 100 106  94 106 100
37  41  50  40  45  68  64  55  69-101-130-103-109-127-127-120-122-140-130-124-117-126-120-100 100 102 106  96 101 106
34  45  45  36  41  64  64  45  64  83 -88 -82 -93 -94-107-115-117-106-108-128-122 -94  96 107 101  98  94 106 106  96
34  33  33  34  37  40  41  41  49  55  64  69  68  85  92  89  89 101 100  91  83  83  93  78  83 106  91  75  91 107
27  34  34  35  36  35  40  41  41  41  45  55  55  64  67  68  75  67  70  69  75  64  67  83  76  67  76  82  83  70
24  25  22  38  26  30  34  36  36  40  37  37  42  66  50  49  55  65  61  58  55  67  73  64  61  65  58  64  67  73
33  34  26  33  27  36  26  27  34  38  36  35  36  45  41  37  40  45  42  38  49  54  49  40  42  51  50  41  52  51
```

Figure 23.   NVLI/6--Tank (10x20/137)
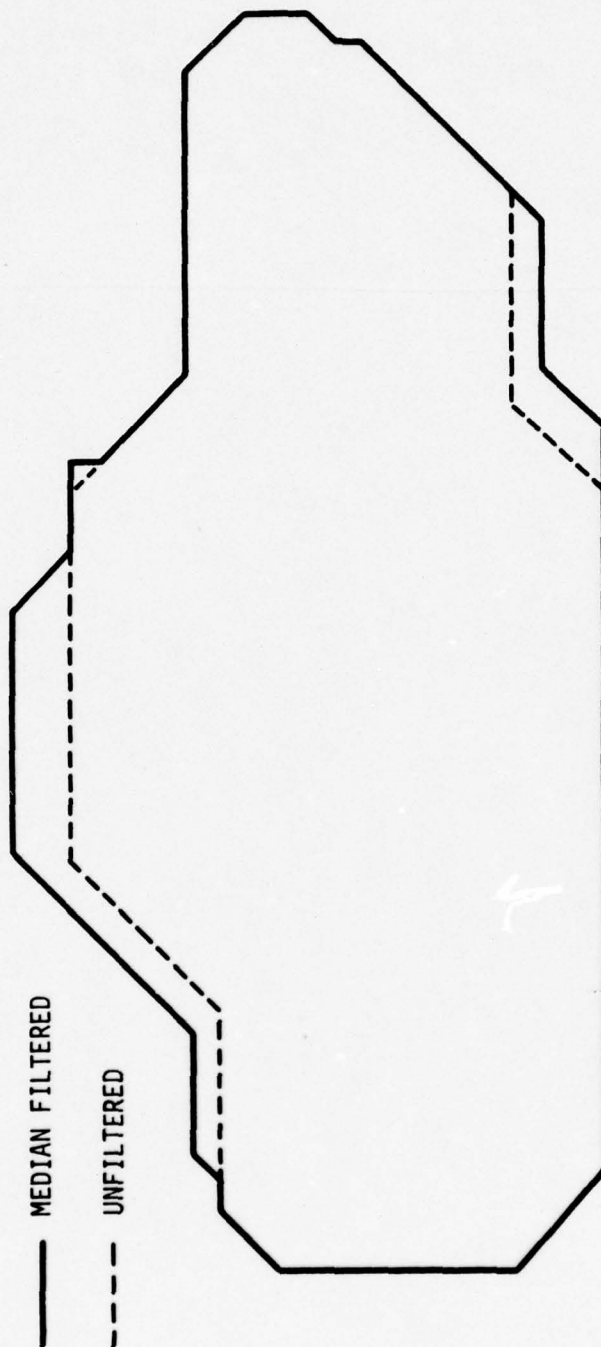
MEDIAN FILTERED

UNFILTERED

Figure 24. NVLI/7--Tank (18x42/550)

Figure 25 delineates the types of statistics which can be extracted from the boundary chain code. The chain code is edited to smooth rough boundaries, and partitioned into edges based upon an average straightness criterion. Size independent statistics such as the ratio of perimeter to root-area ($P/\sqrt{A}$), number of edges, distribution of normalized edge lengths (length/perimeter), and differential slope (change in slope from edge to edge as the boundary is traced in a clockwise direction) are most significant for comparisons amongst target classes over an ensemble of images. Note that $P/\sqrt{A} \simeq 4$ for the 6 tanks is typical of square objects, as opposed to more or less circular blobs where $P/\sqrt{A} \simeq 2\sqrt{\pi}$ would be expected.

Figures 26 and 27 are sample histograms of average intensity difference (contrast) and gradient of the boundary obtained from the boundary extraction routine. It is anticipated that these two features may provide useful information in classifying target characteristics in the future.

BRIGHTNESS STATISTICS

With the target separated from the background, intensity histograms for the target (Figure 28), and local background (Figure 29) are easily extracted. From these, histogram moments and contrast measures such as:

$$\text{Average Contrast:} \quad C_A = \frac{|M_T - M_B|}{\max(M_T, M_B)}$$
and
$$\text{Peak Contrast:} \quad C_P = \frac{|I_{T,\max} - M_B|}{\max(I_{T,\max}, M_B)}$$

can be computed. Here, T refers to the target, B to the background, M is a mean, and $I_{T,\max}$ is the peak target intensity.

57

A cluster diagram of brightness mean versus standard deviation for eleven targets (mixed tanks, jeeps and APCs) is shown in Figure 30. For this limited set of images, targets are eminently separable from backgrounds with the first two moments of the intensity histogram.

HEIGHT:	13	AREA:	275
WIDTH:	27	P/ A:	4.2
PERIMETER:	68	# EDGES:	8.5

NORMALIZED EDGE LENGTH
$m = 0.12$, $\sigma = 0.11$

DIFFERENTIAL SLOPE
$m = -44^{\circ}$, $\sigma = 28.6^{\circ}$
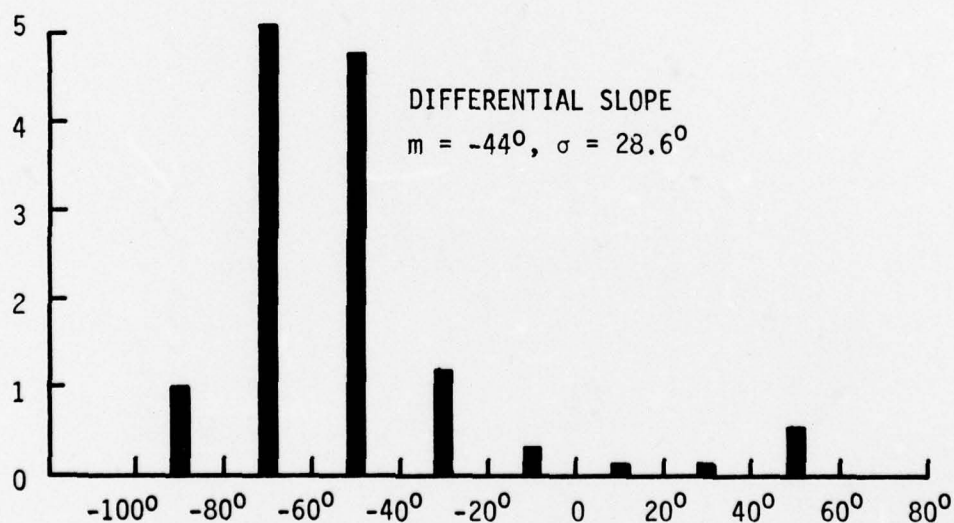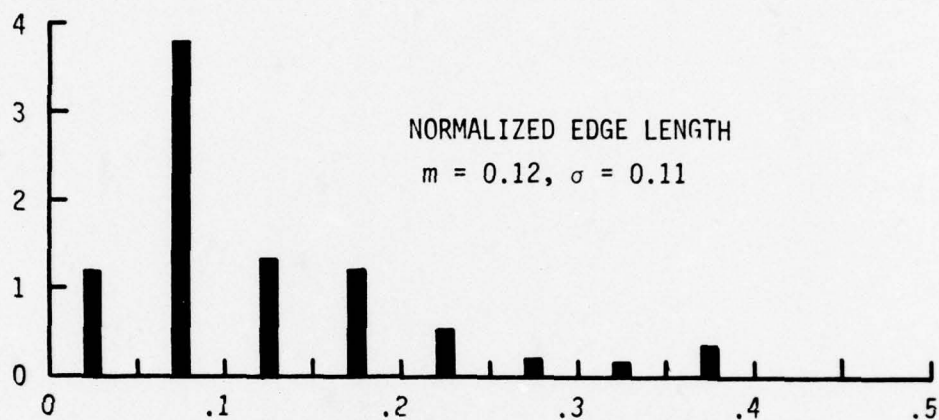
Figure 25. Target Edge Statistics
(Average--6 Tanks)

58
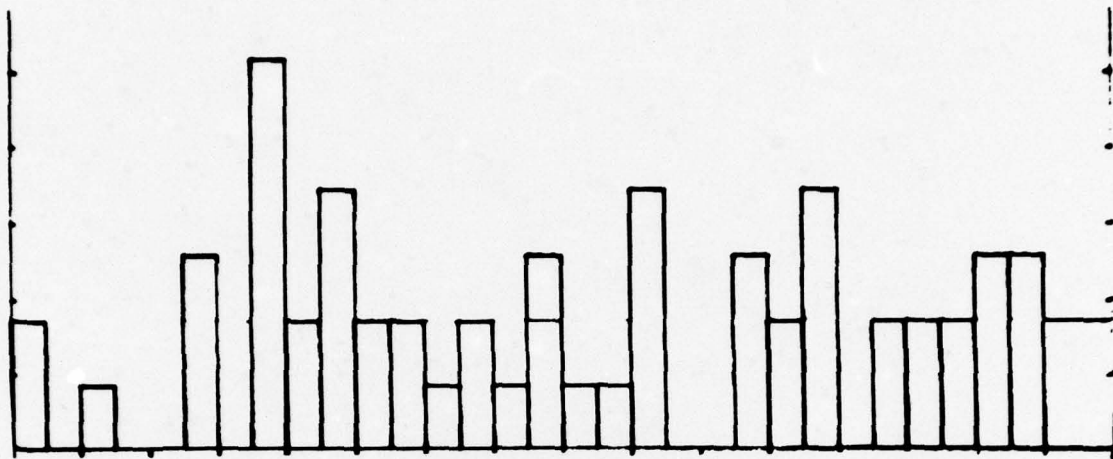
Figure 26.   Average Contrast Sample Histogram
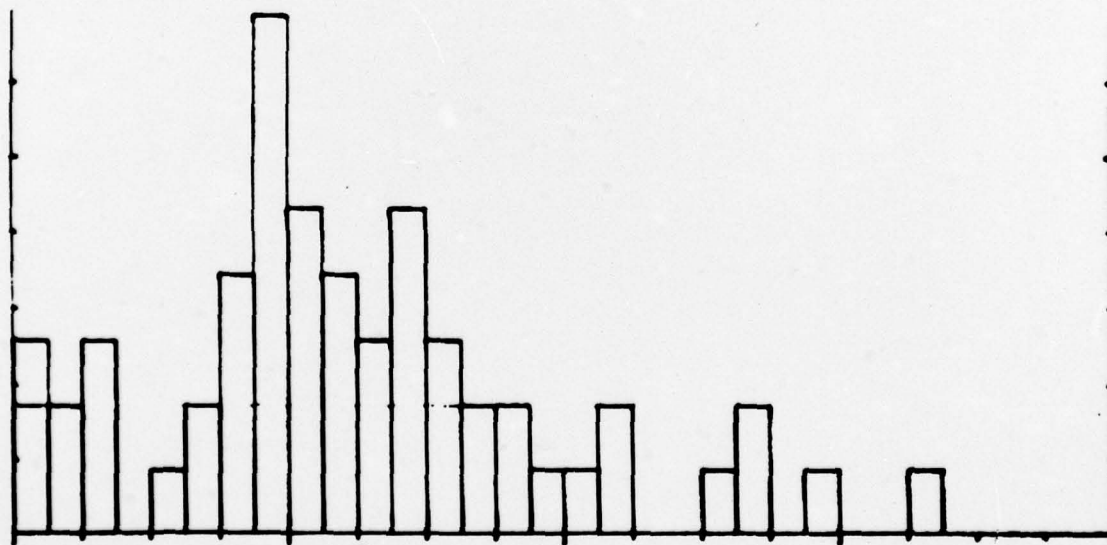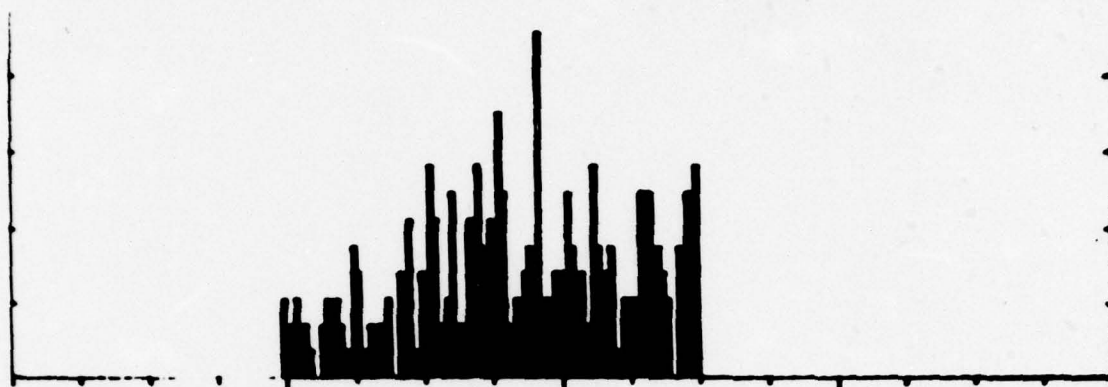

Figure 27.   Boundary Gradient Sample Histogram

59

Figure 28. Target Intensity Histogram

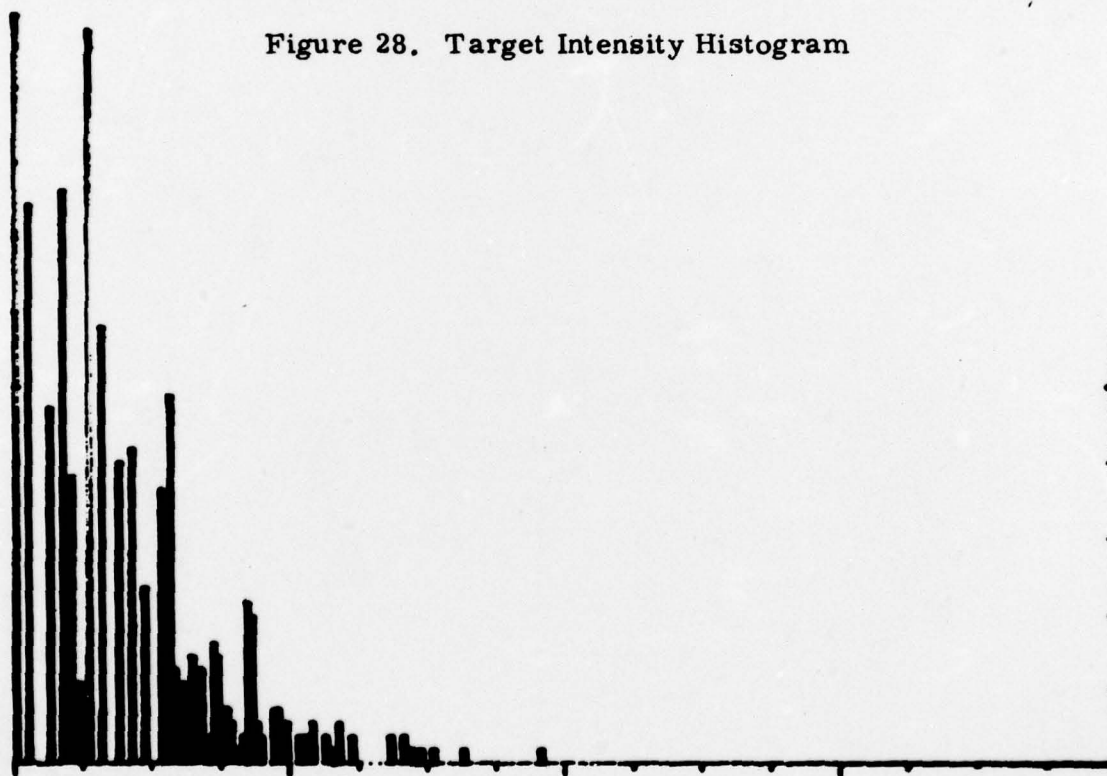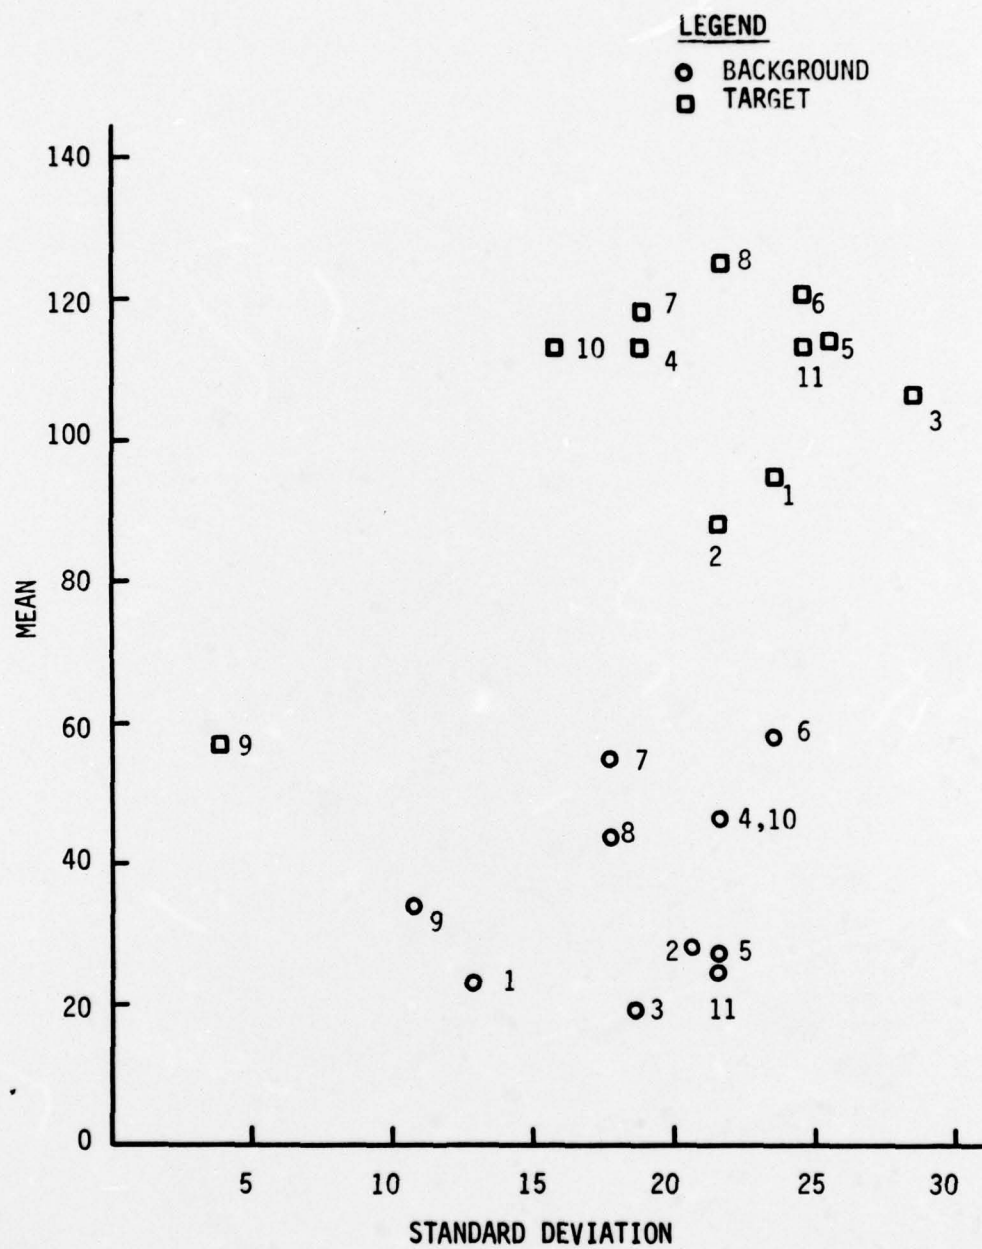Figure 29. Background Intensity Histogram

60

LEGEND
○ BACKGROUND
□ TARGET

140

120

100

MEAN

80

60

40

20

0

□ 8
□ 7
□ 10 □ 4
□ 6
□ 5
11
□ 3
□ 1
□ 2

□ 9
○ 6
○ 7
○ 4,10
○ 8
○ 9
○ 2 ○ 5
○ 1
○ 3 11

5   10   15   20   25   30

STANDARD DEVIATION

Figure 30.   Mean vs. Standard Deviation of Brightness Histograms

61

TEXTURE FEATURES

These are measures to quantify the grey level variations in the image, called texture. We have implemented two classes of texture definitions to evaluate FLIR imagery:

- Grey level difference statistics [Rosenfeld (reference 1)] based on the Haralick (reference 2) measured.

- Texture features based on Max-Min measures developed by Professor O. R. Mitchell of Purdue University.
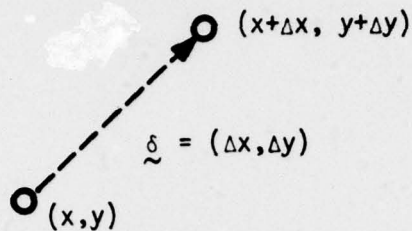
Grey Level Difference Statistics

Assume that the texture is to be measured over a local area of the image. Consider all pairs of points in the region, exactly at a vector distance $\underset{\sim}{\delta} = (\Delta x, \Delta y)$ apart (as in Figure 31). Let $|p(x+\Delta x, y+\Delta y) - p(x,y)|$ be the grey level difference for each such pair. The histogram of these grey level differences (of all pairs of points exactly $\underset{\sim}{\delta}$ apart) is called the grey level difference histogram $p_\delta(\ )$. If there are N grey levels, then there are N bins in this difference histogram. This is a measure of the probability density of grey level differences occurring in the image at a given spacing and orientation. The shape of this histogram is a measure of the texture. The various descriptors (describing this histogram) are:

a. $\frac{1}{N} \Sigma\, i\, p_\delta(i)$ (Mean)

b. $\sum\limits_{i=0}^{N-1} i^2\, p_\delta(i)$ (Contrast)

62

$$\underset{\sim}{\delta} = (\Delta x, \Delta y)$$

$$\underset{\sim}{\delta} = \{(0,1),\ (1,0),\ (1,1),\ (1,-1)\}$$
$$\{(0,2),\ (2,0),\ (2,2),\ (2,-2)\} \quad . \ . \ .$$

Figure 31. Illustrating Point Pairs for Grey Level Difference Statistics

c. $\quad \displaystyle\sum_{i=0}^{N-1} p_\delta^2(i)$ (Angular Second Moment)

d. $\quad -\displaystyle\sum_{i=0}^{N-1} p_\delta(i)\ Ln[p_\delta(i)]$ (Entropy)

Figure 32 gives the grey level difference histograms for a target and its immediate background area for $\delta = (0,1),\ (2,2),\ (4,0)$ and $(8,8)$.

The difference histograms offer a great deal of information on the texture. But we need fewer descriptors of texture than the full histograms. Hence the above measures (mean, contrast, etc.) were computed from these histograms. These were computed on the target (tanks) and background areas for six NVL supplied images; for $\underset{\sim}{\delta} = [(0, \Delta),\ (\Delta, 0),\ (\Delta, \Delta),\ (\Delta, -\Delta)]$ corresponding

TARGET   BACKGROUND

$\delta = (0, 1)$

$\delta \approx (2, 2)$

$\delta \approx (4, 0)$

$\delta = (8, 8)$
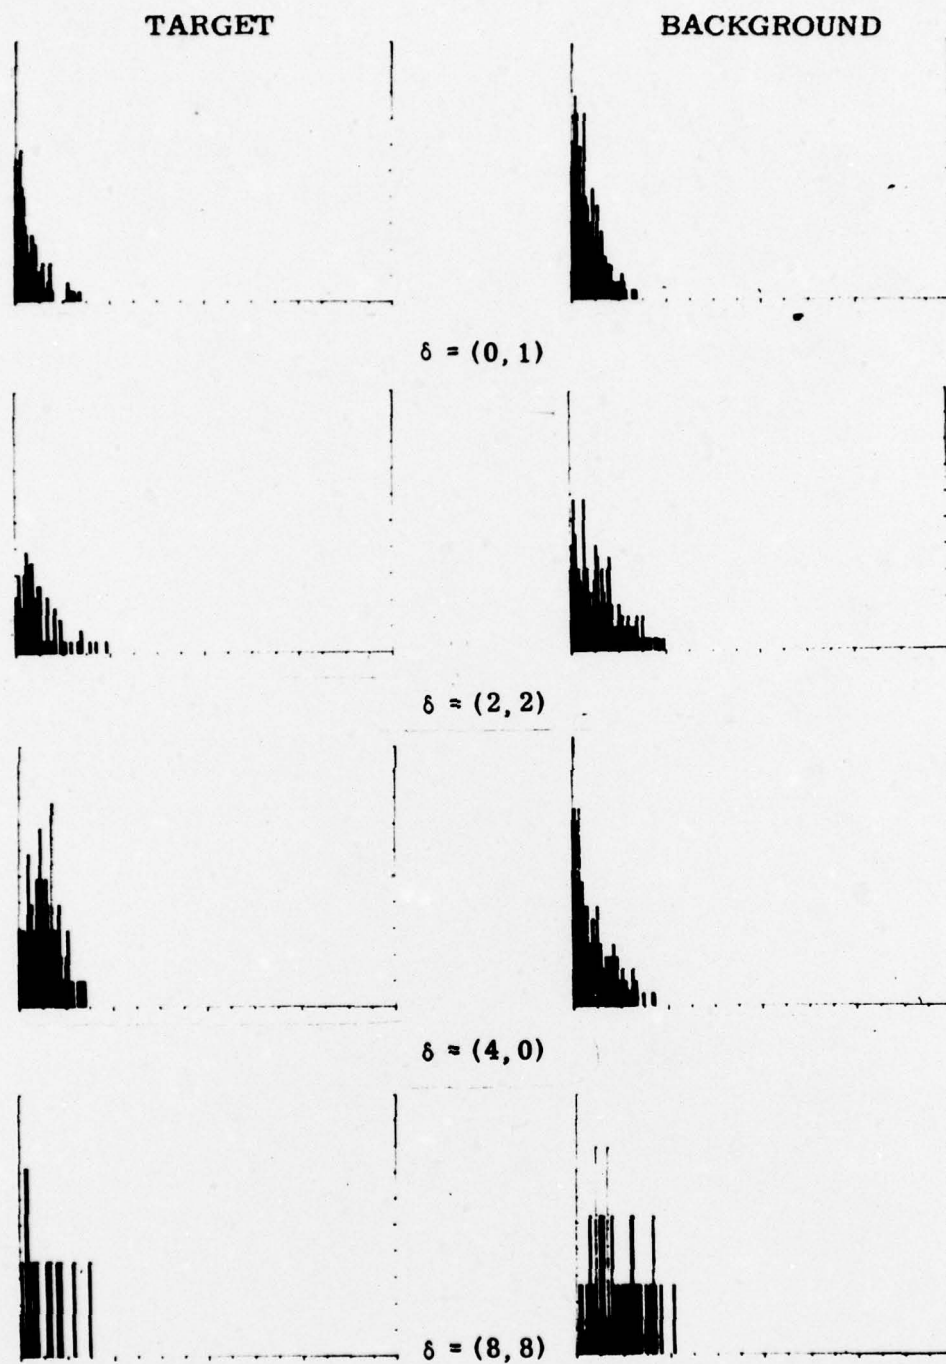
Figure 32.   Grey Level Difference Histograms for a Representative Target
and Its Background for Various Values of δ

64

to orientations of 0°, 90°, 45° and -45° and spacings of Δ = 1, 2, 4 and 8. The parameters (mean, contrast) were averaged over all directions for a given spacing, Δ. Figure 33 is the corresponding plot of the mean measure for the various values of Δ = 1, 2, 4 and 8 for the six tank targets and their respective backgrounds. Note that the separation in this feature for each target and its background increases with Δ as might be expected. But unfortunately the measures do not cluster into background and target values. This might also be expected because the background can possess a variety of textures. However, this is too small a sample to be statistically meaningful. Figure 34 is a similar plot for the contrast measure.
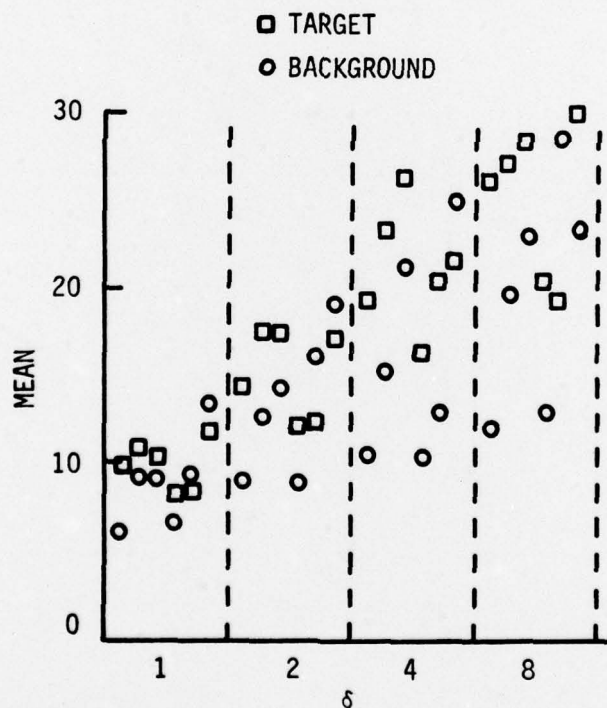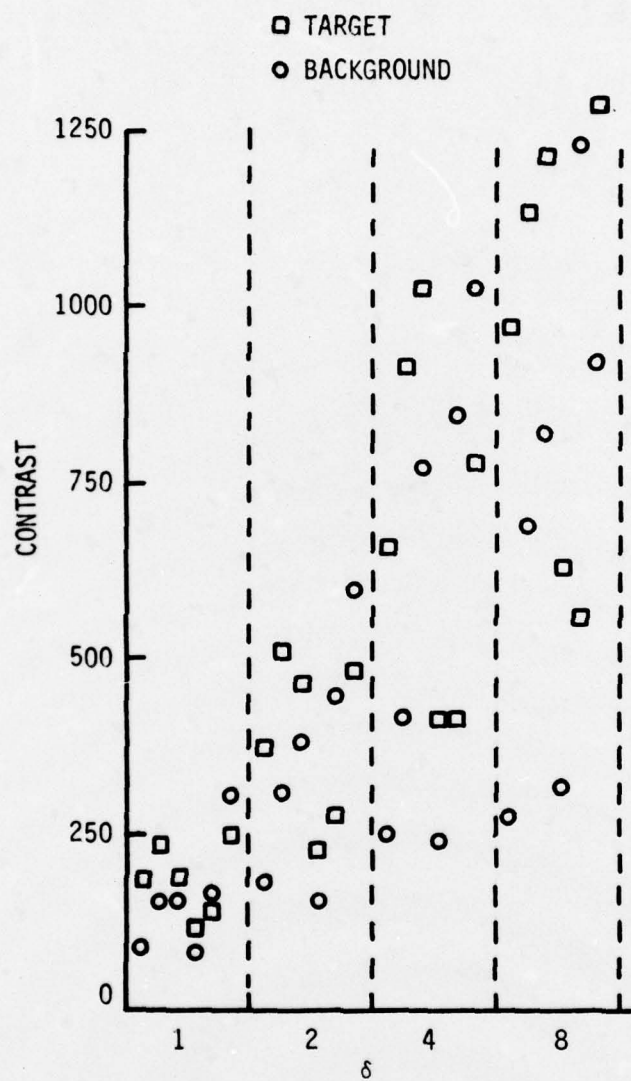

Figure 33. The Mean Measure for Various Separations δ

65

Figure 34.   Contrast Measures for Various δ

Max-Min Texture Descriptors

The general idea here is that the number of extrema (i. e., maxima and minima) per unit length along a given direction in the image is a measure of the image texture. Since we are dealing with irregular areas in the image (such as target/background) this concept can be generalized. The distances between successive extrema are measured in a given direction in the area under consideration. Figure 35 illustrates this. The distances are measured and histogrammed along the rows and columns separately. The average distance, the variance of these distances, as well as skewness and excess of these histograms are computed from the distance histograms. The rows and columns are further subjected to hysteresis filtering with different lags (see the section on MRT enhancement); the extrema are computed and the distances measured for each hysteresis lag. Hysteresis filtering with different lags yields extrema of different magnitudes. The corresponding distances can be heuristically related to spatial frequency concepts, with 1/(mean distance) corresponding to the average frequency.

Figure 36 shows the histogrammed distances between extrema for target and background areas for an NVL image for different hysteresis lags, 0, 4, 8 and 12.

The mean and standard deviation of these distances (measured along rows) for each of 11 targets and its backgrounds are plotted in Figures 37 and 38 for lags 4 and 8 respectively. Again, note that the target and backgrounds do not cluster in the texture measures, but each target differs in texture from its background.

67

Although the texture measures presented above do not appear to cluster into background and target values directly, they do yield discriminatory information between the target and its background. It will be interesting to see how these texture features are transformed by the contrast enhancement and resolution restoration processes.
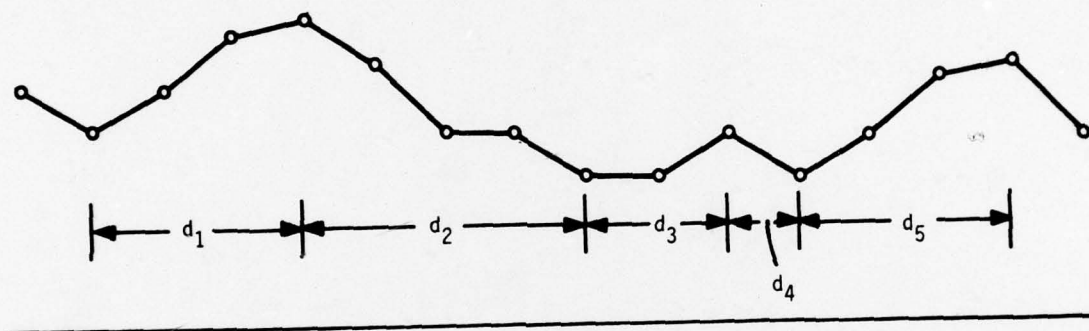
Figure 35. Illustrated Distances Between Extrema Along a Given Direction in the Image
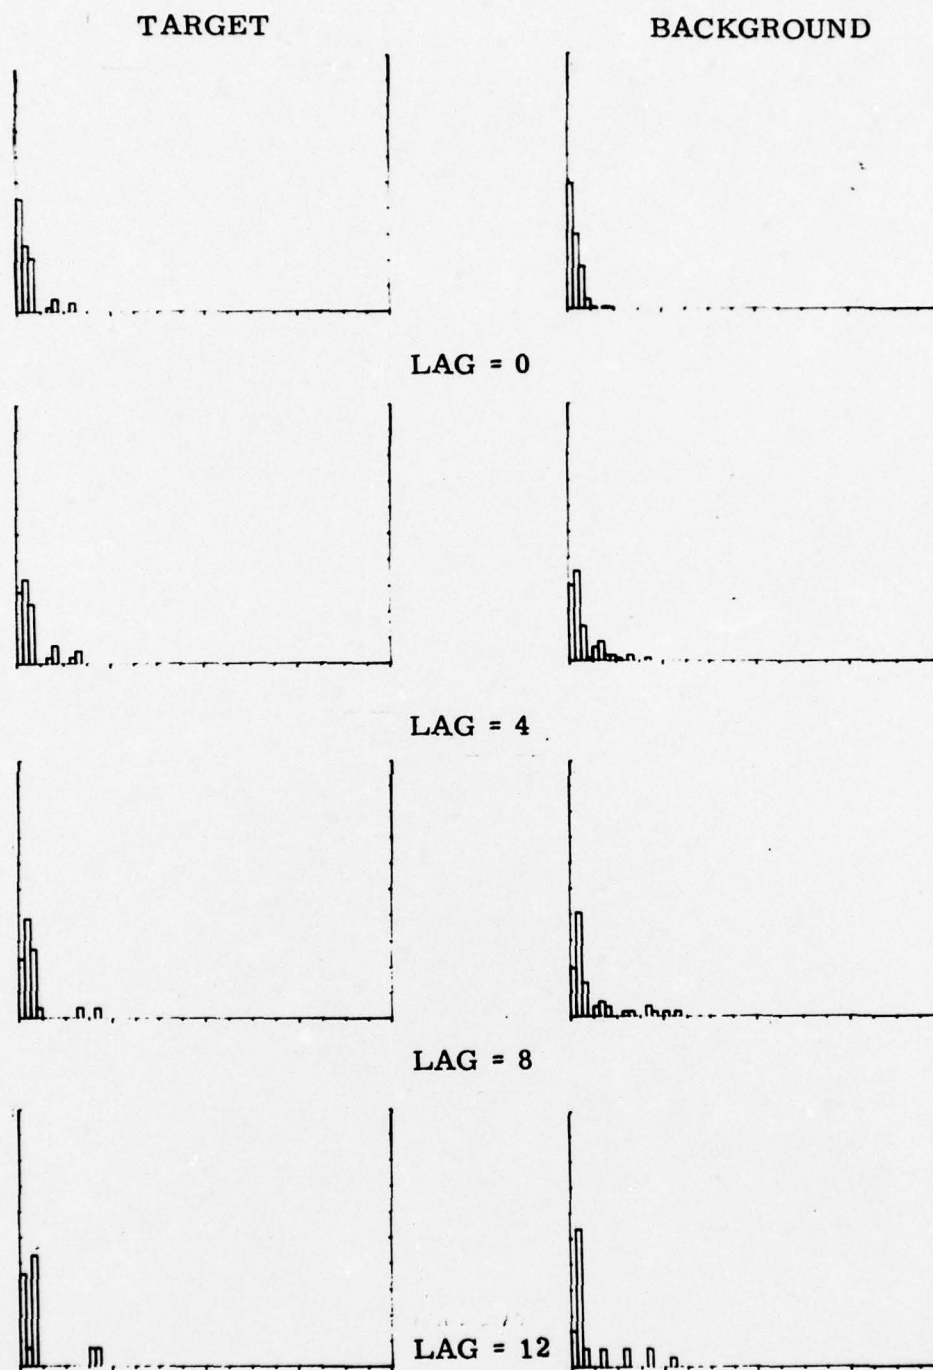
TARGET   BACKGROUND

LAG = 0

LAG = 4

LAG = 8

LAG = 12

Figure 36.   Histograms of Distances Between Extrema for a Target
             and Its Background, for Different Hysteresis Lags 0, 4,
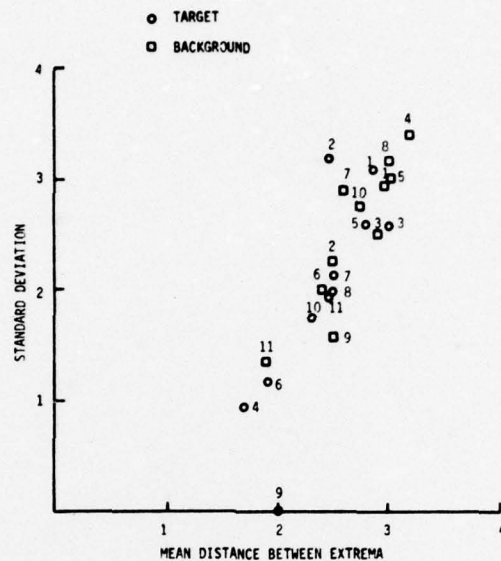   8 and 12

69

Figure 37. Plot of Mean and Standard Deviation of Distances Between Extrema for Hysteresis Lag = 4
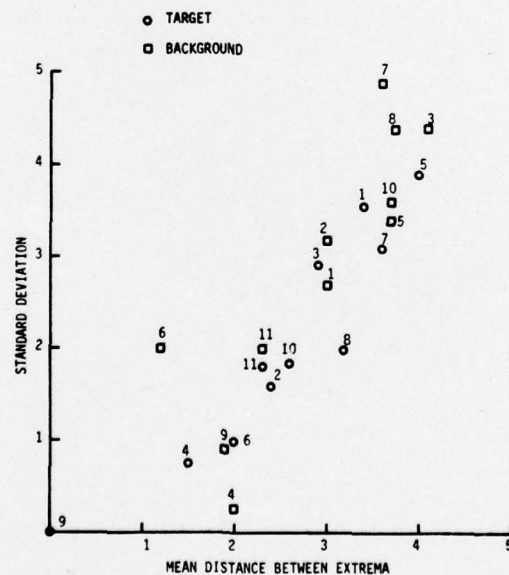

Figure 38. Plot of Mean and Standard Deviation of Distances Between Extrema for Hysteresis Lag = 8

70

REFERENCES

1. Wezka, J. S., Dyer, C. R. and A. Rosenfeld, "A Comparative Study of Texture Measures for Terrain Classification," IEEE Trans. Systems, Man, and Cybernetics, Vol. 6, April 1976, pp. 269-285.

2. Haralick, R. M. et al., "Textural Features for Image Classification," IEEE Trans. Systems, Man, and Cybernetics, Vol. 3, November 1973, pp. 610-621.